

Hedge funds, signaling, and optimal lockups*

Juhani T. Linnainmaa

Alan Moreira

November 2017

Abstract

Many hedge funds restrict investors' ability to redeem their investments. We show that lockups alleviate a delegation friction. In our model hedge funds can trade a long-term arbitrage opportunity; doing so increases expected returns but lowers short-term returns. Investors who rationally learn from returns may mistake a skilled manager who pursues the arbitrage opportunity for an unskilled manager. Skilled managers therefore have an incentive to avoid redemptions by distorting their portfolios to enhance short-term returns. The tradeoff between the ability to trade the arbitrage opportunity more aggressively and investors' fears of being stuck with an unskilled manager determines the optimal lockup. We calibrate the model to hedge fund data and show that arbitrage remains limited even with optimal lockups; the average manager sacrifices 78 basis points in expected returns per year to improve short-term returns.

*Linnainmaa is with the University of Southern California and NBER. Moreira is with the University of Rochester. We thank Douglas Diamond, Lars Hansen, Zhiguo He, John Heaton, Ralph Koijen, Robert Vishny, Tobias Moskowitz, Pietro Veronesi, Alexi Savov, Michael Gofman, Asaf Manela, Jaroslav Borovicka, Mathias Kronlund, Diogo Palhares, Benjamin Pugsley, Christopher Hrdlicka, Savina Rizova, and seminar participants at University of Chicago, University of California at Los Angeles, City University of Hong Kong, and Yale for helpful comments.

1 Introduction

Many hedge funds impose an initial lockup period to restrict investors' ability to redeem their investments. Funds with lockup restrictions significantly outperform open-ended funds.¹ This performance difference implies that managers also benefit from the self-imposed lockup restriction—if not, managers could enhance their own compensation by removing the restriction. Restrictions on redemptions also take many other forms such as infrequent redemption windows, notification periods, discretionary gates, and back-load fees. Managers employ all these tools to increase capital stickiness. In this paper we model lockups to examine what purpose these tools serve. We show that lockups alleviate a delegation friction that emerges when investors are uncertain about managerial skill.

Our model has the following elements and sequence of events. Investors have prior beliefs about a manager who is about to open a fund. The manager offers investors a contract with a fee schedule and a lockup provision, obtains assets, and trades until the fund is liquidated. Skilled managers earn abnormal returns by selecting securities and by timing an arbitrage opportunity. Security selection enhances short-term returns and the long-term arbitrage opportunity increases long-term expected returns at the expense of short-term returns. Skilled managers know how to trade the arbitrage opportunity to maximize expected returns. Unskilled managers destroy value.

Fund investors learn about managerial skill from realized returns, and they know that skilled managers profit both by selecting securities and by trading the long-term arbitrage opportunity. They cannot, however, directly observe managers' positions nor the arbitrage trade, and so they sometimes abandon skilled, but unlucky, managers. We solve this model for the manager's optimal trading rule, investors' investment and liquidation policies, and lockup maturity.

¹See Aragon (2007) and Agarwal, Daniel, and Naik (2009).

At the heart of the model is a tradeoff that determines the optimal lockup maturity. Investors perceive the lockup as costly when a manager's reputation is poor: they may discover that the manager is unskilled, in which case they are stuck with him for the duration of the contract. Managers, on the other hand, prefer long contracts. With investors locked in, they can maximize expected returns without worrying about the negative signal carried by lower short-term returns. In equilibrium, managers offer, and investors accept, lockup contracts because the manager can transfer a part of the expected gain to the investors as a lockup premium.

Uncertainty about skill is an important feature of the model. If investors know that their manager is skilled, the lockup provision is redundant. But with uncertainty, it gives managers more time to demonstrate that they are skilled. Investors learn about managerial skill from returns, and optimally decide when to withdraw money. This threat of liquidation alters manager behavior. A skilled manager knows that he might be unlucky—an aggressive bet on the arbitrage opportunity might not pay off—in which case investors infer that he is unskilled and withdraw their investments. The manager therefore has the incentive to improve his short-term returns by trading the arbitrage opportunity less aggressively. Managers close to liquidation optimally build reputation by enhancing short-term returns at the expense of expected long-term returns.

Learning and signaling generate a strong amplification loop. A manager who is concerned about being liquidated enhances his short-term returns by trading the arbitrage opportunity less aggressively. Because investors understand that this is the optimal response by a skilled manager, they know short-term performance is now more informative about skill. Investors' increased focus on short-term performance, however, gives the manager an even stronger incentive to focus on short-term performance. The manager's trading horizon thereby shortens even further.

The threat of liquidation acts as a limit to arbitrage. Managers trade the arbitrage opportunity less

aggressively than they would if capital was sticky. In the model, the extent to which skilled managers leave expected returns on the table measures the limits to arbitrage. For example, a manager might be able to obtain an expected return of 10% per year if he did not need to worry about signaling; but in equilibrium, with the liquidation threat, his optimal portfolio might earn an expected return of just 6%. In this example, signaling concerns limit arbitrage by 4%.

Investors find long-term contracts less attractive also because they face competition from other investors. A manager who turns out to be skilled attracts more capital. This competition drives rents of the manager's human capital towards him. If, on the other hand, he turns out to be unskilled, investors are stuck with him for the duration of the contract. Competition for skill therefore introduces a costly asymmetry. Investors dislike long contracts because they are more likely to bind when the manager is unskilled.

We calibrate the model to match the salient features of the hedge fund industry. These features include the lockup premium of 4%, fees, the attrition rate of young hedge funds, and the distribution of the lockup maturities. We show that, without lockup contracts, the average new manager forgoes 167 basis points in expected returns to build up reputation. Optimal lockups decrease this distortion to 78 basis points. These distortions weaken over time as managers build reputation. Among those who survive a year, the average distortion falls to 29 basis points.

Our paper relates to three strands of literature. First, our model builds on Shleifer and Vishny's (1997) insight that arbitrage may be limited because investors may infer from poor short-term performance that an arbitrageur is not as competent as they thought, and withdraw capital. In our model, both investors and managers are rational in their delegation and investment decisions, and in how investors learn from past returns, and yet the Shleifer and Vishny (1997) mechanism significantly curtails arbitrage in equilibrium.

Second, we complement the literature in which investors learn about managers' abilities from returns. In models such as those by Berk and Green (2004), Dangl, Wu, and Zechner (2006), Basak, Pavlova, and Shapiro (2007), and Berk and Stanton (2007), investors rationally learn from past returns but fund managers do not face a signaling problem that would distort their portfolio choices. In our model, investors also infer skill from past returns, but managers moreover alter their behavior to signal ability.

Third, the contracting and signaling problem resembles models in which corporations and financial intermediaries use the maturity of their liabilities as a signal. Diamond (1991) and Stein (2005), for example, study an environment in which good managers signal their type by choosing short-term contracts. In our model, the underlying signaling problem is different, and the optimal contract maturity depends on how costly entrenchment is and how much a long-term contract increases expected returns by facilitating arbitrage.

Our model is a stylized description of the hedge fund industry. It incorporates only the elements that are necessary for capturing the key tradeoff: investors dislike long-term contracts because the manager might be unskilled, and skilled managers prefer these contracts to be able to trade the arbitrage opportunity aggressively without worrying about short-term signaling effects. Even this simple setup, however, gives rich intuition for the forces that determine the optimal contract maturity. Lockups may also serve other purposes. For example, a mechanism complementary to ours relates to illiquidity. If a fund holds illiquid assets, it may restrict redemptions to mitigate trading costs. In addition to allowing managers to exploit long-term arbitrage opportunities, managers can use lockups to match the illiquidity of their assets and liabilities (Cherkes, Sagi, and Stanton 2009).

We discuss our model as a description of the hedge fund industry, but its mechanism applies more broadly. It describes any principal-agent setup in which the principal is uncertain about the agent's abilities, and the agent can signal ability by enhancing short-term performance at the expense of the

long term. Our model implies that a long-term contract can alleviate the distortions that emerge from the signaling-liquidation problem, but that significant distortions may remain.

2 Model

The model has the following sequence of events. The manager first chooses a contract that specifies the lockup maturity, and then raises capital from investors. The manager then begins trading, and continues to do so until the fund is liquidated. In addition to the fund's current investors, there are outside investors who search and bid for skilled managers.

2.1 Managers' investment opportunities

A manager is either skilled (denoted by $\theta = g$) or unskilled (denoted by $\theta = b$). A skilled manager generates abnormal returns by selecting securities and by trading a long-term arbitrage opportunity. The manager always generates alpha by selecting securities, but he can choose his exposure to the arbitrage opportunity. The stock-selection strategy is exposed to idiosyncratic risk. The long-term arbitrage opportunity is exposed to "crashes:" it earns a riskless return except when there is a crash, in which case its value jumps. The idiosyncratic risk is represented by standard Brownian motion dB_t ; the crash risk of the arbitrage opportunity is represented by a Poisson process dN_t with an intensity δ_ξ .

A dollar invested by a manager in the security selection strategy evolves as,

$$dS_t = \left(r + \alpha^\theta \right) dt + \sigma dB_t + \tilde{\omega} dN_t, \tag{1}$$

where r is the riskless rate, α^θ is the manager's alpha, σ is the volatility of the security selection strategy in the absence of crashes, and $\tilde{\omega}$ is distributed normally with mean zero and a variance of ω^2 . This

variance represents the riskiness of the security selection strategy during a crash. The skilled manager generates excess returns on his selection portfolio, while the unskilled manager tries to select securities but loses money, $\alpha^g > 0 > \alpha^b$.

The price of the long-term arbitrage opportunity evolves as,

$$dA_t = (r + \lambda) dt - \xi_t dN_t. \quad (2)$$

In normal times, this opportunity returns $\lambda > 0$ in excess of the riskless rate; when it crashes, it returns $-\xi_t$. Because crashes are infrequent, this opportunity generates positive returns most of the time. If the crash is sufficiently small, the expected excess return on this opportunity is positive; if the crash is large, it is negative. We assume that ξ_t takes one of two values, ξ^h or ξ^l . If $\xi_t = \xi^h$, the impending crash is so large that the expected excess return is negative, $E[dA_t - rdt] = \lambda - \xi^h \delta_\xi < 0$. If the impending crash is small, the expected return is positive. That is, the crash sizes satisfy $\xi^h > \lambda/\delta_\xi > \xi^l$.

We assume that the expected return on the long-term arbitrage opportunity is zero; that is, without knowing the size of the impending crash, a manager cannot profit by trading it. The two values of ξ_t are equally likely. The value of ξ_t is fixed until there is a crash, after which its value is redrawn from the unconditional distribution.

The return on the manager's portfolio, which combines the security-selection strategy with the long-term arbitrage opportunity, is

$$dR_t = dS_t + \pi_t(dA_t - rdt). \quad (3)$$

The skilled manager knows whether ξ_t is high or low. He invests $\pi_t \in [-1, 1]$ in the arbitrage opportunity, financing this position by buying or selling the riskless asset. The unskilled manager does

not know whether the impending crash is small or large, and so he always takes a long position to earn the carry λ by choosing $\pi_t = 1$. The unskilled manager's strategy is as risky as that of the skilled manager and so fund investors cannot distinguish between the two from volatilities alone.² Because funds keep their sizes constant at one—we describe their dividend policies below— R_t is the cumulative dollar gross return up to date t .

2.1.1 Discussion

The assumption that the arbitrage opportunity is exposed to crash risk is a modeling device for capturing the idea that the arbitrage trade is a long-term investment opportunity. The unskilled manager cannot profit by trading this opportunity, but the skilled can. When, unbeknownst to investors, the impending crash is small, the manager faces no signaling problem: the trade that maximizes long-term expected returns is the same that maximizes short-term returns. However, when the impending crash is large, the skilled manager must short the arbitrage opportunity and pay λ in order to profit from the eventual crash. This is the important state, because here the manager has to decide between maximizing long-term expected returns and signaling skill by enhancing short-term returns. If the manager maximizes long-term expected returns, investors may interpret his low short-term returns as evidence of lack of skill and redeem their investments.

The arbitrage opportunity is a positive carry strategy. These strategies earn high Sharpe ratios in normal times, but they are prone to crashes.³ Strategies that fit this characterization include, for example, carry trade (Burnside, Eichenbaum, Kleshchelski, and Rebelo 2011), merger arbitrage (Mitchell and Pulvino 2001), and momentum (Jegadeesh and Titman 1993). Hedge funds trade many of these strategies (Cochrane 2011, p. 1087). Agarwal and Naik (2004) show that equity hedge funds'

²If the unskilled manager followed a strategy that was more or less risky than that of the skilled manager, investors would immediately identify him as being unskilled. He must therefore mimic the skilled manager to delay discovery.

³See, for example, Barroso and Santa-Clara (2015) and Moreira and Muir (2016).

payoffs resemble those obtained by writing put options on the market index; this strategy earns a small premium in normal times but bears significant crash risk.

2.2 Contracts

The contract between the manager and the investors specifies the lockup maturity T , management fee f , and a performance fee γ for returns in excess of the riskless rate. When the lockup expires, all investments become liquid; before this date, all capital is locked up.

The lockup expiration is a Poisson event with an intensity $\delta_e = \frac{1}{T}$. Investors who invest in a fund with a T -year lockup therefore expect the lockup to expire in T years. By assuming stochastic expiration, the problem is not time-dependent; that is, the passage of time carries no additional information about the expiration of the lockup. This assumption is appropriate because the long-term arbitrage opportunity is also a Poisson process—that is, its expected payoff also does not vary with the passage of time.

The fund pays dividends to investors so that it always has one unit of capital under management. With a management fee of f and a performance fee of γ , these dividends equal

$$dD_t = rdt + (1 - \gamma)(dR_t - (r + f)dt), \tag{4}$$

where dR_t is the fund return.

2.3 Investors, perceived return dynamics, and the market for skill

Risk-neutral investors discount cash flows at the riskless rate r . Each fund has both existing investors as well as potential investors. The arrival of these other investors, who search for undervalued fund managers, follows a Poisson process. Because this search process captures the amount of competition in the market for skill we denote its intensity δ_c . The higher the intensity, the more likely it is that a

skilled manager attracts new capital. All investors value managers the same way. We first solve for this valuation and then discuss how new investors bid.

A fund's value to investors depends on their belief that the manager is skilled and the fund's lockup status. Investors have a prior belief P_0 that the manager is skilled. Investors update their beliefs P_t as Bayesians by observing returns. In Section 3.1 we describe how investors update their beliefs. They also form beliefs about skilled managers' strategies. For convenience we will use the \hat{Z}_t notation to mean that \hat{Z}_t is the expected value of the Z_t under the investor information set, in contrast to the skilled manager information set Z_t . We denote investors' beliefs about the manager's portfolio choice by $\hat{\pi}_t(\xi, \theta)$, which depends on the investors' time- t information set $\{P_t, L_t\}$ and also the crash state and manager type. Investors do not know the crash state ξ_t or the manager type θ , but they understand that the manager's optimal choice π_t depends on it; in our notation, they believe that if $\xi_t = \xi^h$ and $\theta = g$, the manager chooses $\hat{\pi}_t(\xi^h, g)$. Because the unskilled manager portfolio is constant and equal to one $\hat{\pi}_t(\xi, b) = 1$, we will omit the type θ and use $\hat{\pi}_t(\xi^h)$ and $\pi_t(\xi^h)$ to refer to the skilled manager portfolio from here on.

Investors value a fund at either $V(P_t, L_t = 1)$ or $V(P_t, L_t = 0)$ depending on its lockup status, where $L_t = 1$ indicates that the lockup is active. If the lockup is active, they value one dollar in the fund at

$$V(P_t, L_t = 1) = \hat{E}_0 \left[\int_0^\tau e^{-rt} dD_t + e^{-r\tau} (1_{\tau_c} + 1_{\tau_e} V(P_{t+\tau}, L_{t+\tau} = 0)) \right], \quad (5)$$

where τ is the time of arrival of outside investors (τ_p) or lockup expiration (τ_e), whichever comes first. If outside investors arrive and bid and the manager accepts the bid, current investors get back their capital. If the lockup expires, investors' new valuation depends on the manager's reputation at the time of this event, taking into account that the fund is now open ($L_{t+\tau} = 0$). Between time t and one of these

events, investors earn dividends from the fund. When the fund's lockup has expired, investors can cash out at will. Investors liquidate if the value of one dollar invested in the fund is less than one dollar, i.e. as soon as the fund value is lower than the fund net asset value. Because fund expected returns are increasing in the manager reputation, this redemption decision is of the threshold type, with investors cashing out as soon as the manager reputation reaches a low enough reputation, or conversely, with investors investing in the fund as soon as the manager has a high enough reputation. We will label this threshold $p_l(L)$, where the argument L makes it clear that this threshold depends on the lockup status of the fund. We discuss these thresholds further in Section 3.

When outside investors notice an undervalued fund, which happens with intensity δ_c , they value the fund as in equation (5). If they value the fund at a premium, they make an offer that transfers some of the surplus to the manager. If they value the fund at a discount, they do not bid,

$$C(P_t) = \begin{cases} \beta [V(P_t, L_t = 1) - 1] & \text{if } V(P_t, L_t = 1) \geq 1, \\ \emptyset & \text{otherwise.} \end{cases} \quad (6)$$

Here, β represents the competitiveness of the bidding process. If $\beta = 1$, bidding is perfectly competitive and the manager receives all the surplus. If the manager accepts the bid, current investors have their capital returned to them, and the new investors take their fund space.

In our baseline model we assume that funds impose lockups on any new investors regardless of the lockup state of the existing investors. This assumption is consistent with how the industry operates, and it is also convenient from the modeling perspective.

2.3.1 Discussion

The arrival of outside investors represents competition for skill. Competition drives rents of the manager's human capital back towards the manager. In Berk and Green (2004), capital markets are competitive, and so investors capture none of these rents. The literature has explored different mechanisms that capture this rent transfer. Sirri and Tufano (1998), for example, assume that successful managers are promoted to manage larger funds. Our assumption that outside investors search for skilled managers is the same mechanism that Berk and Stanton (2007) propose as an explanation for the closed-end fund discount puzzle. The mechanism matters less than the outcome. If there is a mechanism that allows managers to receive pay increases when their perceived value increases (as in Holmström and Harris (1982)), that pay increase always comes at the expense of the existing investors. When outside investors successfully bid for an undervalued manager, current investors receive their capital back—but these are also the states in which they value the manager the highest.

2.4 Manager

Both skilled and unskilled managers discount cash flows at the riskless rate r . By managing a fund, a manager earns the management fee f and performance fee γ specified in the contract. Because the manager knows that he might receive outside offers in the future, his valuation today also reflects such potential offers. The manager's valuation therefore depends on the contract terms, his reputation, and the arrival rate of the outside offers.⁴ The skilled manager's reputation and compensation both depend

⁴Because the unskilled manager is non-strategic—he always trades to earn the carry from the arbitrage opportunity—we only need the skilled manager's value function.

on his portfolio choices, and he chooses an investment policy π_t that maximizes his value:

$$G(P_t, L_t, \xi_t) = \max_{\pi_t} \mathbb{E} \left[\int_0^{\tau} e^{-rt} [\gamma(dR_t - (f+r)dt) + f dt] + 1_{\tau_c} e^{-r\tau_c} [G(P_{t+\tau_c}, L_t = 1, \xi_t) + C(P_{t+\tau_c})] \right. \\ \left. + 1_{\tau_e} e^{-r\tau_e} G(P_{t+\tau_e}, L_{t+\tau_e} = 0, \xi_t) + 1_{\tau_l} e^{-r\tau_l} \underline{G} \right], \quad (7)$$

where \underline{G} is the value of an outside option that the manager receives if the fund is liquidated. The first term represents the fee that the manager receives until one of three events occurs: (1) there is an outside offer at time $t + \tau_c$ that is accepted by the manager; (2) the lockup expires at time $t + \tau_e$ (if the lockup was currently in place); or (3) fund investors withdraw their investments at time $t + \tau_l$ (which can happen only if the lockup has expired). The manager always accepts the outside offer because new investors pay a front load fee $B > 0$ and new funds come locked-up.

Every manager who is about to start a new fund is confident that he is skilled. Managers vary in their initial reputations P_0 . The manager first writes a contract that specifies the lockup maturity T , and offers this contract to investors. He chooses the lockup maturity to maximize his valuation,

$$T^*(P_0) = \arg \max_T \left\{ C(P_0) + \mathbb{E}_0[G(P_0, L_0 = 1)] \right\}, \quad (8) \\ \text{s.t. } C(P_0) \geq 0,$$

where the $C(P_0) \geq 1$ condition ensures that investors break even—that is, that a dollar in the fund is valued at least at a dollar and investors bid for fund access. The bidding process is the same as before; investors bid $C(P_0) = \beta [V(P_0, L_0 = 1) - 1]$, where β determines the fraction of surplus that is transferred to the manager. After starting a fund, a skilled manager makes his decisions according to equation (7) to maximize his value. An unskilled manager always earns the carry from the arbitrage

opportunity—these managers can be viewed as either maintaining their belief that they are skilled (and consistently underestimating the size of the impending crash), or as learning their type after they start the fund, but choosing to earn the carry to improve their short-term returns. Because the unconditional expected excess return on the long-term arbitrage opportunity is zero, the unskilled manager’s portfolio choice does not affect his long-term expected return, while at the same time increasing the fund performance in the short-term.

2.5 Equilibrium

We assume that investors and managers play the equilibrium that maximizes expected returns and, therefore, everyone’s valuations. We solve a Markovian equilibrium where the state variables are the manager’s reputation (P_t), the fund’s lockup status (L_t), and the crash state (ξ_t). Investors’ decisions can only depend on the manager’s reputation and the lockup status; their beliefs about skilled managers’ actions, however, can also depend on the crash state that they do not know. Skilled managers’ decisions depend on their reputation and the current crash state. The equilibrium imposes consistency between the investors’ actions and beliefs, and the actions of skilled managers.

Definition 1. Markovian equilibrium. *A Markovian equilibrium is given by skilled managers’ portfolio π_t^g , investors’ beliefs about this portfolio $\hat{\pi}_t^g(\xi)$, a law of motion for investors’ beliefs P_t , and liquidation policies $l(P_t)$, such that:*

- (1) *Skilled managers’ portfolio choices are optimal given investors’ beliefs, liquidation policies, and outside offers.*
- (2) *Investors’ beliefs about skilled managers’ portfolio choices are consistent with skilled managers’ portfolio choices.*

- (3) *Investors' beliefs about managerial skill are consistent with Bayes' rule on the equilibrium path.*
- (4) *Current investors' liquidation policies are optimal given competing investors' bidding behavior, and their beliefs about skilled managers' portfolio choices and the manager's likelihood of being skilled.*
- (5) *Skilled managers' portfolio choices maximize expected returns subject to conditions (1) through (4).*

3 Model solution

The solution proceeds in three stages. We first describe how investors update their beliefs. We then discuss the key properties of skilled managers' optimal portfolio choice. Last, we discuss the investors' optimal investment policies. Appendix B contains all propositions and their proofs. The skilled manager's and investors' valuations satisfy a system of coupled differential equations. Appendix C describes how these equations can be solved.

3.1 Learning about managerial skill

Investors update their beliefs that the manager is skilled as Bayesians. The learning problem divides into two parts. In periods without a crash, investors face the problem of differentiating between two processes with different means but identical variances. When a crash occurs, investors update their beliefs differently as they observe a very large one-time return. The extent to which investors learn from the crash depends on their beliefs about the skilled and unskilled managers' portfolio choices. If investors believe that both managers have the same exposure to the crash risk, they will not update their beliefs when there is a crash. If, however, they believe that the managers have different exposures, crashes can be highly informative. We assume that, after a crash happens, investors know whether ξ_t was high or low.

We let p_t denote the log-likelihood ratio that the manager is skilled, $p_t = \log\left(\frac{P_t}{1-P_t}\right)$, and use p_t or P_t as the state variable depending of convenience. Proposition 1 in the appendix shows that the log-likelihood ratio of investors' beliefs about skill evolves as

$$dp_t = \frac{\kappa_t}{\sigma} (dR_t - \check{\mu}_t dt) + \frac{\kappa_t^c(\xi_t)}{\omega} (dR_{t+} - \check{\mu}_t^c(\xi_t)), \quad (9)$$

in which dR_t is the instantaneous non-crash return, dR_{t+} is the crash-event return (if there is a crash at time t), κ_t and $\kappa_t^c(\xi_t)$ are the signal-to-noise ratios of short-term performance and crash performance, $\check{\mu}_t$ represents investors' expectations in log-likelihood space about performance in periods without a crash, and $\check{\mu}_t^c(\xi_t)$ represents investors' expectations (in log-likelihood space) about performance in periods crash performance given a particular crash state realization.⁵ We provide the formulas for these variables in the Appendix.

Equation (9) shows that if investors expect skilled managers to display better short-term performance, they infer from positive return surprises ($dR_t - \check{\mu}_t dt > 0$) that the manager is more likely skilled than not. Investors measure performance relative to how well they expect the average manager to perform, $\hat{\mu}_t$. The informational content of short-term performance is measured by κ_t ; it is the difference between skilled and unskilled managers' expected short-term performance, scaled by return volatility. The informational content of crash performance is measured by $\kappa_t^c(\xi_t)$; if the crash is large relative to ω , investors significantly revise their beliefs if $\hat{\pi}_t^g \neq 1$, as long as investors believe that the skilled manager invested differently from the unskilled manager. A key aspect of the learning problem is that in the absence of a crash investors know neither the state ξ_t nor whether the manager is long or short the arbitrage opportunity. They therefore have to form beliefs about managers' choices, and to learn from performance by averaging across the two crash states.

⁵Expectations in log-likelihood space are the equal-weighted expected performance of each manager type.

3.2 Manager's optimal portfolio choice

We take as given the lockup maturity T and let $G(t) = G(p_t, L_t, \xi_t)$ denote the skilled manager's value function from equation (7) (which is now defined in log-likelihood space p_t). The value process $G(t)$ is a martingale for $t < \tau_l$ that satisfies the following HJB equation:

$$\begin{aligned}
 rG(t) = & f + \gamma(\alpha^g + \pi_t(\lambda - \delta_\xi \xi_t) - f) + \delta_c C(p_t) \\
 & + \delta_c (G(p_t, L_t = 1, \xi_t) - G(t)) + \delta_e (G(p_t, L_t = 0, \xi_t) - G(t)) \\
 & + G_p(t) \frac{\kappa_t}{\sigma} (\mu_t - \check{\mu}_t) + \frac{1}{2} G_{pp}(t) (\kappa_t)^2 + \delta_\xi (\mathbb{E}_t [G(p_{t+}, L_t, \xi_{t+})] - G(t)).
 \end{aligned} \tag{10}$$

When the lockup has expired, the valuation respects the boundary condition imposed by investors' liquidation policy so that $G(p_t, L_t = 0, \xi_t) = \underline{G}$ for all $p_t \leq p_l(L = 0)$, where $p_l(L = 0)$ is the reputation threshold at which investors optimally cash from the fund (more on this threshold later). That is, the manager receives the value of the outside option when investors withdraw their investments.

Equation (10) provides intuition for the forces that shape the manager's decisions. On the left we have the time discounting of the value function. On the right, the first line represents the managed expected compensation: management fee, expected performance fee due to non-crash and crash periods, and load fees due to the arrival of new investors' flows.

In the second line we have the valuation effects of changes in the lockup status of the fund. The first term captures the valuation gain from new investors being initially locked up; it is zero if the lockup is already active for the current investors. The second term measures the valuation loss of a lockup expiration; it is zero for funds with already expired lockups.

In the last line the first two terms capture the valuation effect of expected reputation changes in periods without a crash. The expected growth rate is increasing in the wedge between a manager's

and investors' beliefs about expected fund performance. This growth rate increases or decreases in the volatility of reputation, κ_t/σ , depending on whether the value function is convex or concave in reputation. The third term is the valuation effect of expected reputation changes in periods with a crash. The expected reputation jump in a crash is $E_t[p_{t+} - p_t] = \frac{1}{2}(\kappa_t^c)^2$. A crash has a large impact on the manager's valuation if the signal-to-noise ratio κ_t^c is high and learning is sharp, i.e. if the equilibrium differences across skilled and unskilled portfolios is large relative to the amount of idiosyncratic risk ω to which managers are exposed.

Equation (10) is a system of four coupled integro-differential equations—these are distinct from ordinary differential equations in that they feature non-local movements in reputation. The number of equations is four because the system jumps between the two crash states ($\xi_t = \{\xi^h, \xi^l\}$) and between funds with active and expired lockups ($L_t = \{0, 1\}$).

Propositions 2 and 3 show that the manager's valuation, with some restrictions on the parameters, is increasing and S-shaped in his reputation. This S-shape induces state-dependent variation in the manager's aversion to reputation shocks: a manager with a low reputation prefers reputation shocks while a highly reputable manager is averse towards them. Proposition 4 shows that lockups reduce managers' short-term reputation concerns.

We can examine the determinants of the skilled manager's portfolio choice by focusing on the terms in equation (10) that depend on π_t .

$$\max_{\pi_t} \underbrace{\pi_t \gamma (\lambda - \delta_\xi \xi_t)}_{\text{compensation incentives}} + \underbrace{\pi_t \lambda \frac{\kappa_t}{\sigma} G_p(t)}_{\text{short-term}} + \overbrace{\delta_\xi E_t \left[G \left(p_t + \frac{\kappa_t^c(\xi_t)}{\omega} (-\pi_t \xi_t + \tilde{\omega} - \check{\mu}_t^c(\xi_t)), L_t, \xi_{t+} \right) \right]}^{\text{reputation concerns}} \underbrace{\hspace{10em}}_{\text{long-term}}, \quad (11)$$

where the expectation is with respect to the manager's information set at time t .

Equation (11) shows that the objective is almost linear in the portfolio choice. The only non-linearity stems from the manager's concerns about long-term (that is, post-crash) reputation. Therefore, absent these concerns, the manager's portfolio choice will be at one of the extremes, at either $\pi_t = -1$ or $\pi_t = 1$.

Performance fee γ gives the manager an incentive to maximize expected returns. Other things equal, an increase in this fee leads the manager to place more weight in maximizing expected returns. The performance fee term in equation (11) is positive when the impending crash is small ($\xi_t = \xi^h$) and negative when it is large.

The second term in equation (11) represents short-term reputational concerns. This term is positive when $\kappa_t > 0$, that is, when investors view high short-term returns as a positive signal about managerial skill. This short-term reputational effect gives the manager an incentive to take a long position in the long-term arbitrage opportunity regardless of the size of the impending crash. The more informative short-term returns are about skill, the more the manager distorts his portfolio.

The last term in equation (11) represents long-term reputational concerns. A manager's portfolio choice also affects the manager's reputation once the crash happens. This term balances out some of the short-term concerns: while a manager might be tempted to take a long position in the arbitrage opportunity even when the impending crash is large, he is concerned about what doing so would do to his reputation if there is a crash. These long-term concerns therefore incentivize the manager to maximize expected returns.

The relative importances of the short- and long-term reputational concerns depend on the manager's horizon. This horizon appears in equation (11) as the difference in the slope of the value function with respect to reputation (G_p) before and after the crash. These slopes are similar when a manager has a long horizon; for a manager close to liquidation, however, the before-crash value function is much

steeper.

Proposition 5 describes the manager’s optimal portfolio choice. When the impending crash is small, the manager’s optimal strategy is always $\pi_t = 1$; in this case, the strategy that maximizes expected returns is the same that maximizes short-term returns and therefore best signals skill. When the impending crash is large, the optimal choice ranges from $\pi_t = -1$ to $\pi_t = 1$ depending on the strengths of the reputational effects. If the short-term reputational concerns are small enough, the manager maximizes long-term expected returns by choosing $\pi_t = -1$. If these short-term concerns are large, they can completely overwhelm the long-term reputational concerns. The manager then maximizes short-term returns at the expense of long-term expected returns by choosing $\pi_t = 1$. There is also an intermediate region with $-1 < \pi_t < 1$; here, the marginal long-term reputational concerns exactly balance out the short-term concerns and payoff incentives.

Figure 1 illustrates how the equilibrium portfolio choice depends on reputation. Here, a skilled manager manages a fund without a lockup clause, knowing that the impending crash is large, $\xi_t = \xi^h$. If the manager’s reputation is too low—here, if the probability that the manager is skilled is less than $1/3$ —investors are unwilling to invest with the manager. This is the “liquidation region” in the graph.

If the manager’s reputation is low but above the liquidation threshold, the manager has a strong incentive to maximize short-term performance. He thus invests $\pi_t = 1$ in the long-term arbitrage opportunity despite knowing that he will suffer significant losses if there is a crash. This is the optimal choice because the manager is so close to the liquidation threshold; he cannot afford to maximize long-term expected returns because, in expectation, it will take too long for the arbitrage opportunity to pay off.

As the manager’s reputation improves in Figure 1, the manager’s reputational concerns lessen, and the manager shifts towards the strategy that maximizes long-term expected returns. However, only

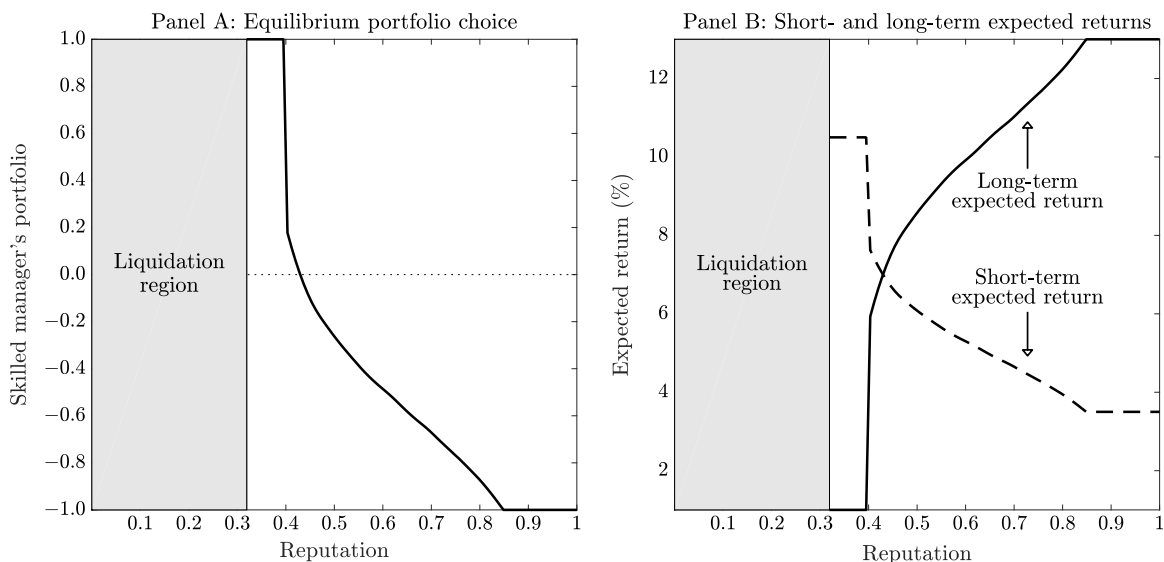


Figure 1: **Skilled manager's equilibrium portfolio choice and signaling.** This figure shows how the skilled manager's equilibrium portfolio depends on his reputation when $\xi_t = \xi^h$. The skilled manager runs a fund without a lockup clause. Panel A shows the optimal portfolio and Panel B shows the expected short-term and long-term returns. The expected short-term return is the expected return conditional on no crashes. The shaded area on the left denotes managers whose reputation is too low to raise capital. As the manager's reputation increases, the manager does not need to improve his short-term returns to signal skill, and therefore invests in the strategy that maximizes expected returns, $\pi_t = -1$.

when the manager's reputation is very high—here, if the probability that the manager is skilled is above 0.8—he follows the strategy that maximizes long-term expected returns, entirely ignoring short-term signaling issues.

Figure 1 shows that reputational concerns limit arbitrage. A skilled manager with a low reputation willingly forgoes expected returns to signal his ability to ensure survival. The gap between the maximum attainable expected return and the expected return that the manager accepts in the equilibrium measures limits to arbitrage. In Figure 1 the maximum expected return, obtained by shorting the arbitrage opportunity is 13%; however, a manager whose reputation is just barely above the liquidation threshold earns an expected return of just 1%. Such a manager trades against the arbitrage opportunity, making a gamble that there will not be a crash until he has built enough reputation to pursue the expected

return-maximizing strategy.

Proposition 6 shows that the limits to arbitrage increase as managers get closer to the liquidation threshold. This result holds for managers who manage funds without lockups and for those whose funds have lockups and sufficient reputation. It also shows that the limits to arbitrage decrease in the length of the lockup.

A long-term contract affects a manager's equilibrium choice by lowering the value he puts on short-term changes in reputation. The lockup provision ensures that the investors cannot liquidate the fund even if the manager's reputation temporarily drops below the level at which investors would like to do so. The lockup provision therefore provides insurance against temporary drops in reputation. It allows the manager to take a long-term view to portfolio choice.

3.2.1 Feedback loop and equilibrium multiplicity

The portfolio choice examined in Equation (11) depends not only on exogenous expected returns $\lambda - \delta_\xi \xi_t$ and endogenous reputation concerns G_p , but also on investors beliefs about the manager portfolio choice through the terms κ_t , κ_t^c , and $\check{\mu}_t^c$. This feature is typical of learning environments where the content of the signal depends on the expected action, and in turn the action taken depends on the signal content. This mechanism leads to amplification and equilibrium multiplicity.

Intuitively, a reduction in reputation that increases reputation concerns G_p increases the incentives to perform well in the short-term and as a result increase π towards one. Investors realize this change in incentives and adjust upwards their beliefs about the gap between the short-term performance of skilled and unskilled managers κ_t . This further increases incentives to perform well in the short term through second term in Equation (11) and increases π further.

Multiplicity arises naturally if reputation concerns G_p are large relative to compensation incentives.

Investors beliefs become self-fulfilling as their learning behavior provide sufficient incentives for the skilled manager to conform with their expectations. Similar considerations arises for the long-term reputation concerns because κ_t^c and $\tilde{\mu}_t^c$ depend on investors and shape manager portfolio incentives.

Because investors and managers can communicate we focus on the equilibrium with the highest expected returns (and therefore lowest amount of distortions), as this equilibrium features higher compensation for the manager and higher net distributions to investors.

3.3 Investors' optimal policy

Investors' value function satisfies a HJB equation similar to that of the manager. The important difference between the two is the return dynamics. Whereas the skilled manager knows these dynamics, under the investors' information set they are a mixture of the skilled and unskilled managers' dynamics. The mixture weight is the manager's current reputation. Investors' HJB equation can be written as,

$$\begin{aligned}
rV(t) = & [r + (1 - \gamma)(\hat{\mu}_t + \delta_\xi \hat{\mu}_t^c - f)] + \delta_c(1 - V(t)) + \delta_e [V(p_t, L_t = 0) - V(t)] \\
& + \frac{1}{2} V_{pp}(t) \kappa_t^2 + \delta_\xi \left(\hat{E}_t[V(p_{t+}, L_t)] - V(t) \right), \tag{12}
\end{aligned}$$

where the expectation integrate out investors' uncertainty about the manager type θ , the crash state ξ_t , and the idiosyncratic risk realization $\tilde{\omega}$. The investors' valuation depends on the fund expected net of fees return (first term in the right), the intensity of competition for managers (second term), the expected time to lockup expiration (third term), and how quickly investors expect to learn about managerial skill (second line).

Investors' optimal policy consists of choosing the smallest manager reputation at which they are willing to invest when the shares are liquid, and the smallest reputation at which they are willing to invest

when the shares are illiquid for a maturity T . These thresholds are the solution of $V(p(0), L = 0) = 1$, $V_p(p(0), L = 0) = 0$, and $V(p(1), L = 1) = 1$. The first and third equality are “value matching” conditions that state that the investor must be indifferent between one dollar inside or outside the fund at the liquidation threshold. The second equality is a “smooth pasting” condition that arises because when the fund is open ended the investor has the option to delay divestment. This flexibility implies that the value function must have zero slope at the optimum threshold. There is no smooth pasting condition for the optimal investment in the fund because investors cannot postpone investment in the fund. Once they find a fund, they must invest or lose the investment opportunity.

Proposition 7 shows that the liquidation threshold is increasing in the arrival rate of other investors. That is, when investors are concerned that good news about the manager’s reputation will be promptly captured by outside investors, they are less willing to invest with a manager with a low reputation. The same proposition also shows that when the market for skill is sufficiently competitive, the liquidation threshold increases in the limits to arbitrage. That is, investors are unwilling to invest in low-reputation managers when they know that these managers’ equilibrium response is to enhance short-term returns at the expense of expected returns.

4 Implications of the model

In this section we calibrate the model to the hedge fund industry. We use this calibration to illustrate why and when skilled managers distort their portfolio choices, and how these distortions change as managers build reputation. We also show how the lockup contracts alleviate managers’ reputational concerns, and discuss the factors that determine the optimal lockup maturity.

Table 1: Calibrating model to hedge fund data

This table reports the parameter estimates that we use to calibrate the model to hedge fund data. The last column, when applicable, references the study that estimates the parameter, or states how the parameter is selected.

Parameter	Notation	Value	Note
Risk-free rate	r	1%	
Fund fees			
Performance fee	γ	20%	Agarwal et al. (2009)
Management fee	f	1%	Agarwal et al. (2009)
Fund risks			
Idiosyncratic fund volatility	σ	11%	Agarwal et al. (2009)
Crash volatility	ω	5.5%	
Long-term arbitrage opportunity			
Carry	λ	4%	Jurek and Stafford (2015)
Crash intensity	$\delta\xi$	0.33	
High crash	ξ^h	24%	Calibrated to give skilled managers up to 6% returns from market timing
Low crash	ξ^l	-3%	
Security selection abilities			
Skilled manager's alpha	α^g	6%	Match the average lockup premium of 4%
Unskilled manager's alpha	α^b	-5%	
Outside offers			
Arrival rate	δ_c	4	Berk and Green (2004) and Aragon (2007)
Rent sharing	β	1	

4.1 Calibration to hedge fund data

Table 1 summarizes the model parameters and gives the calibrated values; it also references, when applicable, the study that estimates each parameter. These are the baseline values. We later demonstrate how, for example, the signaling problem and the optimal lockup maturity respond to changes in these parameters.

We choose the management fee ($f = 1\%$), performance fee ($k = 20\%$), and idiosyncratic volatility ($\sigma = 11\%$) to match the median fees and risk estimates of Agarwal, Daniel, and Naik (2009, p. 2231). We assume that a crash in the model approximates a shock that unfolds over a period of three months, and set the crash volatility to the volatility over a three-month period, $\omega = 11\%\sqrt{3/12} = 5.5\%$.

We choose the carry of the long-term arbitrage opportunity (λ) to match the return of a put-option writing strategy. Jurek and Stafford (2015) show that a strategy that writes out-of-the-money put options matches the key features of hedge fund returns, and they estimate that this strategy earns a CAPM alpha of 7% per year. Mitchell and Pulvino (2001) show that the merger arbitrage strategy generates a total return of 13%, and that an implicit put writing strategy contributes 3.5% to this total. We set the carry equal to the lower end of these estimates, $\lambda = 3.5\%$.

We assume that the expected return on the long-term arbitrage opportunity is zero; that is, managers who do not know whether the impending crash is large or small cannot profit from it. The difference $\xi^h - \xi^l$ together with the frequency of crashes δ_ξ determines the returns available to skilled managers from market timing. We first set $\delta_\xi = \frac{1}{3}$; this choice implies that a large crash occurs once every six years. By then setting $\xi^h = 24\%$, a skilled manager can earn $\frac{1}{2}\delta_\xi(\xi^h - \xi^l) = 6\%$ per year by timing the market.⁶ This is a modest crash in a risk-neutral world. Jurek and Stafford (2015) estimate that the average hedge fund earned a return of -20% in 2008.

The arrival rate of outside investors controls the speed at which managers capture the rents of positive news about their skill. Aragon (2007) and Agarwal, Daniel, and Naik (2009) show that the average fund without a lockup provision delivers a zero net alpha. Such estimates are consistent with a competitive market for managerial skill (Berk and Green 2004). Therefore, to reflect the typical evaluation frequency of hedge funds, we set the arrival rate of outside investors to $\delta_p = 4$. This choice implies that, on average, positive news are diluted in a quarter. We set $\beta = 1$, which is the parameter that determines how managers and investors split rents. In the model, this parameter and the arrival rate of outside offers enter multiplicatively, and so we can vary either β or the arrival rate to alter the amount of competition for skill.

⁶Because the two crash states are equally likely and the expected excess return on the long-term arbitrage opportunity is zero, equation (2) implies that $\frac{1}{2}(\xi^h + \xi^l) = \frac{\lambda}{\delta_\xi}$. The choices of ξ^h , λ , and δ_ξ therefore determine ξ^l .

Parameters α^g and α^b represent skilled and unskilled managers' security selection skills. Two equilibrium conditions restrict these choices. First, unskilled managers have to underperform the risk-free rate, which is the investors' best outside option. Second, the gap between the skilled and unskilled managers has to be sufficiently large so that investors interpret abnormal positive performance as good news about type. For this condition to hold both when the impending crash is large or small, the difference in alphas has to be at least as large as the carry, $\alpha^g - \alpha^b \geq \lambda$. With these restrictions in mind, we use the lockup premium to choose α^g and α^b . Aragon (2007) shows that funds with lockup provisions earn net alphas between 4% and 7% relative to funds with no restrictions; Agarwal, Daniel, and Naik (2009) estimate a 3% to 4% premium. The lockup premium increases in the gap between α^g and α^b because it is about the risk of being tied up with an unskilled manager who earns low returns. We set $\alpha^g = 6\%$ and $\alpha^b = -5\%$ which, in the model, generate a lockup premium of 3.8%.

4.2 Reputation, survival, and optimal portfolio choice

Managers who are concerned about their reputation trade the long-term arbitrage opportunity cautiously; by distorting their portfolios, they can improve their chances of survival. Figure 2 shows how managers enhance their chances of survival. We give managers of different reputations open-ended funds and measure their one-year survival rates. In this computation, all managers are skilled, and therefore equally deserving of surviving. Not everyone, however, survives; no matter what portfolio they choose, they get liquidated if they are unlucky. Because investors learn from returns, they sometimes mistake an unlucky skilled manager for an unskilled manager. This risk increases as we get closer to the liquidation boundary.

The thick solid line in Figure 2 shows the survival rates of managers who make the equilibrium choices. This choice maximizes the manager's own valuation, which depends crucially on survival and

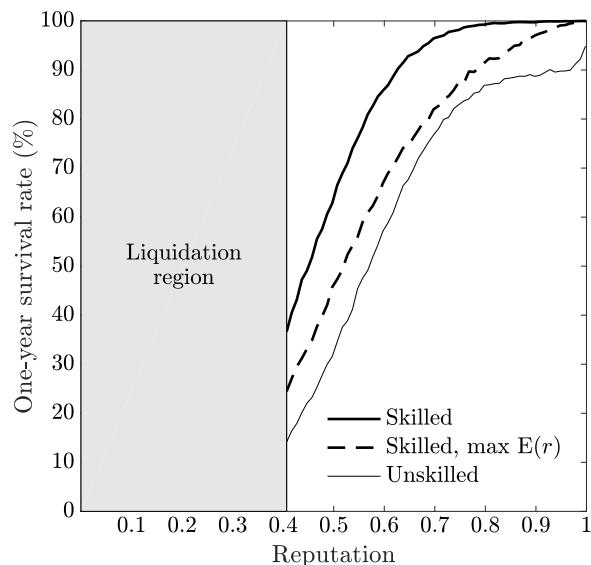


Figure 2: **Survival of skilled and unskilled managers.** This figure shows the one-year survival rates for skilled and unskilled managers of different initial reputations. The thick solid line represents skilled managers who make equilibrium portfolio choices. The thick dashed line represents skilled managers who deviate from the equilibrium by maximizing expected returns. The thin solid line represents unskilled managers who always maximize short-term returns by choosing $\pi_t = 1$.

the expected evolution of reputation. The thick dashed line has the same managers instead choosing portfolios that maximize long-term expected returns. That is, the managers now choose their portfolios without worrying about the signal that their lower short-term returns send to the investors. Here, the manager deviates unilaterally; that is, investors assume that both sets of managers make equilibrium choices. We also report, for reference, the survival rates of unskilled managers. Because unskilled managers, on average, generate lower returns, their survival rates are significantly lower than those of equally reputable skilled managers.

Figure 2 shows that managers' equilibrium choices significantly enhance their survival probabilities. The effect is particularly strong among managers who are close to the liquidation threshold. As shown in Figure 1, these are also the managers who distort their behavior the most. They understand that the long-term arbitrage opportunity typically does not pay off in the short run, and so they instead




Figure 3: **Reputation and limits to arbitrage over time.** This figure simulates five years of data for skilled and unskilled managers. Each manager has an initial reputation of 0.64. Panel A plots the survival rates of these managers. The thick line represents skilled managers who make equilibrium portfolio choices. The thin solid line represents unskilled managers who always maximize short-term returns by choosing $\pi_t = 1$. Panel B plots the average reputation of a skilled manager conditional on surviving. Panel C plots limits to arbitrage for the skilled manager. Limits to arbitrage is the difference between the maximum expected return that the manager could obtain and the one he obtains in equilibrium.

enhance short-term returns to build reputation. As we increase reputations, the gap between the two sets of skilled managers disappears; when a manager’s reputation is high enough, he does not need to distort his behavior to ensure survival.

Figure 3 shows how managers alter their behavior as they build reputation. In this analysis, we examine a manager with a reputation of 0.64. This manager distorts his portfolio moderately to signal skill. Instead of choosing $\pi_t = -1$ to maximize long-term expected returns when the impending crash is large, he chooses a portfolio of $\pi_t = -0.74$.

We now track this manager over time. Even when following the equilibrium strategy, 9 percent of these managers get liquidated over the first year. We also report the survival rates of unskilled managers. The survival rates show that investors therefore typically reach the right conclusion. However, because returns are noisy, some unskilled managers survive for “too long,” and some skilled managers get liquidated despite being skilled.

Panel B of Figure 3 shows that, conditional on survival, skilled managers' reputations improve. If the manager survives for a year, his reputation has increased from 0.64 to 0.74. This increase is due to a combination of two factors. First, the manager's reputation increases because the manager is skilled and earns high returns, and the investors therefore, on average, correctly revise their estimates upwards. Second, because we condition on survival, we cut out the left tail of the distribution, that is, we remove the skilled-but-unlucky managers whose reputations fall below the liquidation threshold.

Panel C of Figure 3 shows the average limits to arbitrage, which is the gap between the maximum available expected return and the equilibrium expected return measures limits to arbitrage. This gap measures the amount of returns that the manager leaves on the table because the reputational cost from harvesting these returns is too high.

The average skilled manager begins to trade the long-term arbitrage opportunity more aggressively as his reputation improves. This is the optimal choice; the manager is no longer as concerned about survival, and so the gain from higher returns offsets the reputational cost of lower short-term returns. In Panel C, however, limits to arbitrage initially increase because of Jensen's inequality. Although the average skilled manager's reputation improves, some experience a decline because of bad luck. Because these unlucky managers tilt their strategies sharply away from the arbitrage opportunity (see Figure 1), the average limits to arbitrage increases. Initially, the average skilled manager foregoes 125 basis points in expected returns to signal skill. This amount increases to 147 basis points in a year before starting the decline. By the end of year two, the average manager gives up just 85 basis points in returns.

4.3 Signaling and commitment

In the model skilled managers sometimes deliberately destroy value to signal skill. Instead of earning an expected return of 10%, a skilled manager may earn a return of just 2% in equilibrium; the threat

of liquidation induces him to distort his behavior. Skilled managers would extract more value from the market if they could commit not to signal; that is, if investors could write and enforce a contract that forces skilled managers to maximize expected returns.

Because investors cannot write such contracts, the set of managers to whom they are willing to provide capital is smaller. As we move below the equilibrium liquidation threshold, the likelihood that a manager is skilled is lower and, if these managers were given assets to manage, they would distort their portfolios the most. We can quantify the effects of the inability to coordinate by changing the model so that, first, skilled managers always maximize expected returns and, second, investors know that skilled managers behave this way and update their beliefs consistent with this knowledge.

In Figure 1, the model parameters are such that the liquidation threshold is $X_L = 0.34$, and the managers close to this threshold maximally distort their portfolios to ensure survival. If investors and managers could coordinate, and the skilled managers would therefore commit to maximizing expected returns, this liquidation threshold would fall to 0.27. That is, investors would now be willing to give capital to managers whose probability of being skilled is just one-quarter. This change in the liquidation threshold can also be expressed as a change in the required return; the higher a manager's reputation, the higher the investors' expectation of his gross return. The change in the liquidation threshold from 0.27 (with commitment) to 0.34 (without commitment) corresponds to investors only considering managers who can deliver 1.4% higher expected gross returns.

4.4 Optimal signaling when short- and long-term returns are more or less informative

Skilled managers distort their portfolio choices because returns carry information that investors use to update their beliefs. Investors learn from both short-term and long-term returns. The informativeness

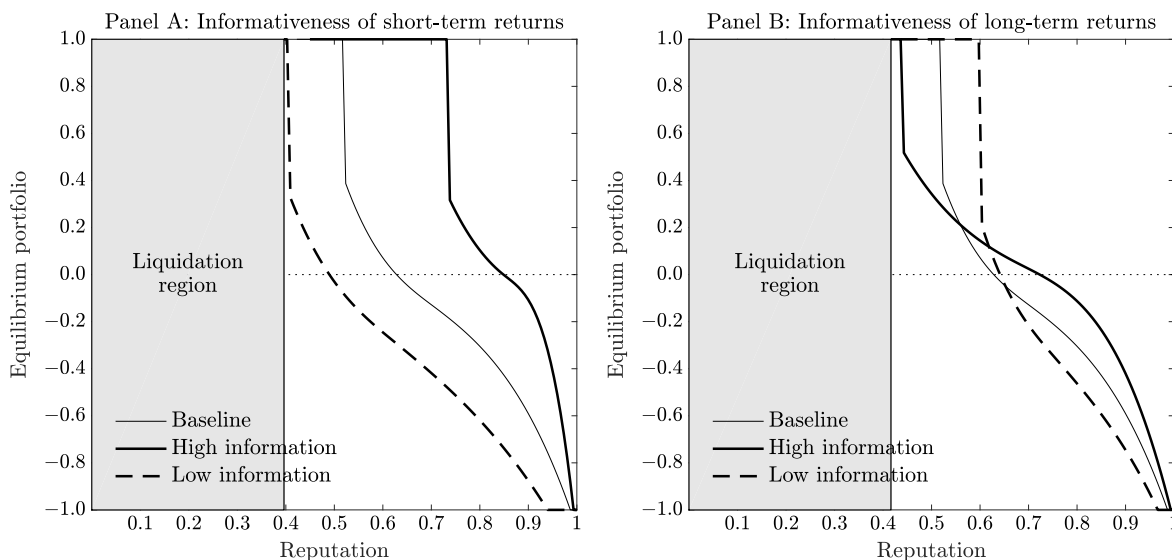


Figure 4: **Optimal portfolio choices when short- and long-term returns are more or less informative.** This figure shows how the skilled manager’s equilibrium portfolio depends on his reputation when $\xi_t = \xi^h$. The skilled manager runs a fund without a lockup clause. Panel A varies the informativeness of short-run returns. The volatility of the Brownian motion, σ , takes the values of 7% (high informativeness), 9% (baseline), and 11% (low informativeness). Panel B varies the informativeness of long-run returns. The volatility of the crash-event returns, ω , takes the values of 9% (high informativeness), 12% (baseline), and 15% (low informativeness). The liquidation boundary also depends on σ and ω ; this figure shows only the lowest thresholds.

of these returns therefore determines the extent to which managers alter their behavior. In Figure 4, we illustrate how the informativeness of short- and long-term returns affect managers’ optimal portfolio choices and therefore limits to arbitrage. Parameters σ and ω determine the speed at which investors learn about managerial skill from short-term and crash-event returns, respectively.

Panel A of Figure 4 plots the equilibrium portfolio choices of skilled managers of different reputations. Similar to Figure 1, we study the choice when the impending crash is large so that the portfolio that maximizes long-term returns invests $\pi_t = -1$ in the arbitrage opportunity. We vary the informativeness of returns around the baseline values. Panel A shows that when short-term returns become more informative (σ falls), managers’ incentives to enhance short-term performance increase. Short-term returns now carry more information. In the baseline case, the threshold below which managers maximally

distort their choices (and choose $\pi_t = 1$) is 0.51. When short-term returns become more informative, this threshold increases to 0.66. All managers below this level view the short-term reputational cost of trading the arbitrage opportunity too high.

Panel B modifies investors' ability to learn from crashes. In Section 2, we considered the two extreme cases where learning during crashes is either perfect or fully absent. At these extremes, manager faces almost no long-term reputational concerns. When learning is perfect, the manager knows that as long as his expected crash performance is slightly better than that of the unskilled manager, investors will learn that he is skilled. Although this result is a consequence of the ability to perfectly detect a jump in a continuous-time environment, it makes the general point that when learning is "lumpy," managers' marginal long-term reputational concerns are weak. At the other extreme, when crash returns carry little information, the manager's portfolio choice today does not affect his expected after-crash reputation. Long-term reputational concerns therefore matter only between these two extremes.

Panel B of Figure 4 illustrates the relations between reputation, the informativeness of crash-event returns, and the limits to arbitrage. When a manager's reputation is high, an increase in the informativeness of crash-event returns increases the limits to arbitrage. However, among managers with low reputation, an increase in informativeness lowers them. That is, when crashes carry more information about skill, low-reputation managers are less eager to distort their portfolios.

4.5 The benefit of lockups to skilled managers

Skilled managers benefit from the lockups because they allow managers to take the long-term view. Managers do not have to worry about short-term returns because, while the lockup is active, investors cannot withdraw their funds even if the manager's reputation temporarily falls under the liquidation threshold. Moreover, if the long-term arbitrage opportunity pays off while the lockup is active, the

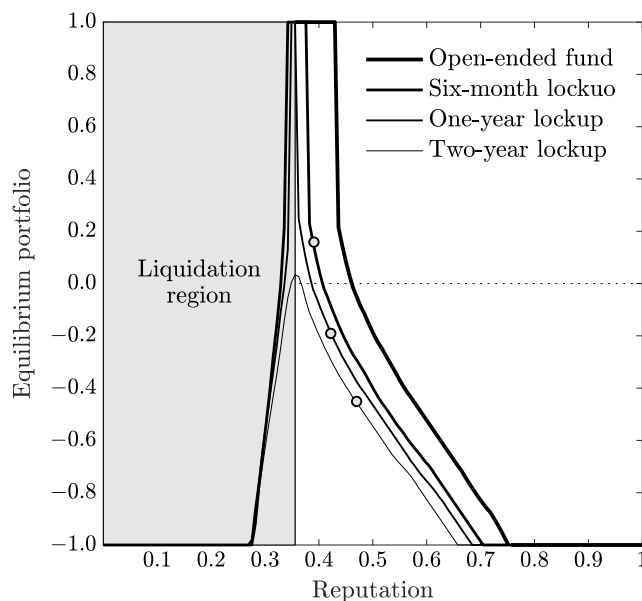


Figure 5: **Optimal portfolio choices with lockup provisions.** This figure shows how a skilled manager’s optimal portfolio choice depends on his reputation when the impending crash is large, $\xi_t = \xi^h$. The manager runs either an open-ended fund or a fund with a six-month, one-year, or two-year lockup. The shaded area signifies the liquidation threshold of an open-ended fund; the circles denote liquidation thresholds of the funds with lockup provisions. Because a manager’s reputation can fall below the liquidation threshold while the lockup is in effect, the three lines that correspond to funds with lockups extend into the liquidation regions.

manager’s reputation will almost certainly exceed that what it would have been had the manager been forced to signal skill by enhancing short-term returns at the expense of long-term returns.

Figure 5 illustrates the benefit of lockups. In this figure, skilled managers of different reputations manage either open-ended funds or funds with lockup provisions ranging from six months to two years. The liquidation threshold also depends on the length of the lockup. The gray area denotes the threshold of an open-ended fund; the gray circles denote the thresholds of the funds with lockup provisions. A manager’s reputations can lie below the liquidation threshold when the lockup is active; if the lockup expires while the reputation is below the threshold, investors withdraw their funds.

Consistent with the intuition that lockups alleviate managers’ short-term reputational concerns, an increase in the length of the lockup increases managers’ willingness to trade the long-term arbitrage

opportunity. When a manager's reputation is very high, the lockup provision does not play any role. Investors are confident that the manager is skilled and, therefore, even a manager of an open-ended fund trades the long-term arbitrage opportunity without worrying about the reputational cost of the low short-term returns. As reputations fall, however, managers reduce their positions in the long-term arbitrage opportunity. Managers who run funds with lockup provisions alter their behavior less because their liquidation risk is lower.

As the reputation of the manager who runs the open-ended fund falls close to the liquidation threshold, he invests $\pi_t = 1$ to maximize his short-term returns. That is, he sacrifices all of the long-term returns to build reputation. The portfolio choices of the managers who run funds with lockup provisions begin to fall back down; that is, they again begin to trade the long-term arbitrage opportunity aggressively. The reason is not that the liquidation risk is low; it is that the liquidation risk is now so high that these managers do not care about any short-term performance boost. The only way these managers can survive after the lockup expires is by trading the arbitrage opportunity and hoping that it pays off in time.

4.6 Optimal lockup maturity

Panel A of Figure 6 shows the optimal lockup maturity as a function of manager's reputation, and Panel B shows how much the manager gains from choosing the optimal lockup over an open-ended contract. When a manager's reputation is close to the liquidation threshold, investors are unwilling to enter anything but an open-ended contract. As the manager's reputation improves, the lockup provision becomes very valuable. These managers are still close to the liquidation threshold, and so long lockups benefit them the most by decreasing the liquidation risk. At the liquidation threshold of 0.41, the optimal lockup is an open-ended contract. However, a manager whose reputation lies 0.05 higher at

0.46 gains 17% in valuation by being able to offer a contract with just a four-month lockup.

As a manager's reputation increases, liquidation risk decreases, and so does the valuation gain. These managers are far away from the liquidation threshold, and so they would benefit less from the lockup provision. Managers who have very high reputations would benefit little from the lockups. However, without additional restrictions, other things equal, they would still weakly prefer long-term contracts. We therefore assume that investors have demand an illiquidity premium, that is, they need to be compensated for tying up their capital for longer periods of time.⁷ Because high-reputation managers are almost indifferent between all contract maturities, even a small cost pushes them to favor short-term contracts.

Panels C and D show how the optimal lockup maturity responds to changes in the manager's market-timing skill and the competitiveness of the market for skill. The amount of returns skilled managers can generate by timing the arbitrage opportunity is an important driver of the optimal maturity. In the baseline calibration, managers can generate an additional 6% in returns by timing the long-term arbitrage opportunity. If skilled managers lose this ability, the signaling problem disappears, and the lockup provision is no longer useful for this purpose. Panel C shows that a manager who chooses a one-year contract in the baseline case with 6% market-timing skill reduces the length of the contract to 4 months when he has no market-timing ability. If, on the other hand, we increase the returns available through long-term arbitrage opportunity to 10%, this manager's optimal maturity increases to 17 months.

Panel D shows that the amount of competition in the market for skill also has a substantial effect on the optimal contracts. An increase in competition makes long-term contracts more costly to investors by

⁷We assume that investors demand a higher hurdle rate in equation (5). That is, instead of discounting at a rate r , they discount at a rate of $r(T)$ that increases in T . This assumption is equivalent to assuming that investors face random liquidity shocks that force them to liquidate funds early and at a discount.

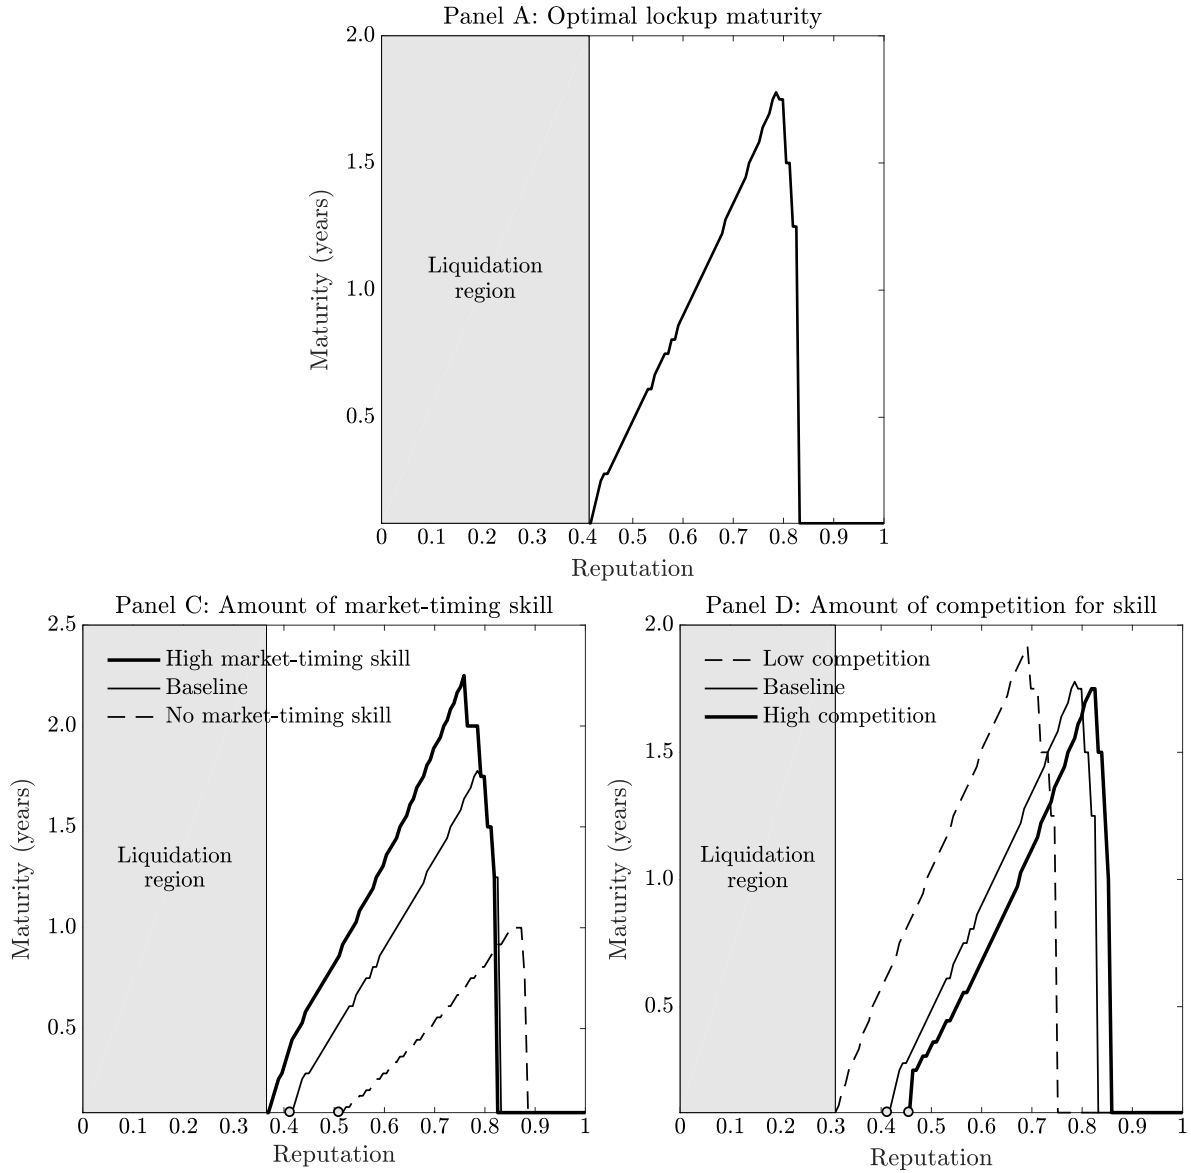


Figure 6: **Optimal lockup maturity: Valuation gains, market-timing skill, and competition in the market for skill.** This figure shows the optimal lockup (Panels A, C, and D) and the valuation gain (Panel D) as a function of manager’s reputation. A manager chooses the optimal lockup to maximize his valuation at the time he sets up a fund. Valuation gain in Panel B is the percentage increase in the manager’s value from moving from an open-ended fund to a fund with the optimal lockup. Panel C varies the amount of market-timing skill; the values are 0% (no market-timing skill), 6% (baseline) and 10% (high market-timing skill). Panel D varies the competitiveness of the market for skill by changing the arrival rate of outside investors. The arrival rates are once a year (low competition), once a quarter (baseline), and once a month (high competition). In Panels C and D, the liquidation threshold depends on the parameters. The shaded area signifies the lowest liquidation threshold; the circles denote the liquidation thresholds of the other cases.

strengthening the asymmetry: skilled managers are discovered more quickly, and so an increase in the arrival rate of outside investors makes it more likely that long-term contracts bind only when investors wish they did not. Panel D shows that when the frequency of outside offers decreases from once a quarter to once a year, a manager who previously chose a one-year contract now writes a 19-month contract.

5 Implied distribution of reputation

Optimal lockup maturities and survival rates depend on managers' reputations. Managers with very high or low reputations will write short contracts, and managers close to the liquidation threshold have lower survival rates than those far above the threshold. In this section, we use these relations between reputations, optimal contracts, and survival rates to back out the implied distribution of reputation.

We assume that managers' (log-likelihoods of) reputations evolve according to equation (??) even before they try to open a fund. Managers' initial reputations are therefore normally distributed. We denote this distribution by $N(\mu_s, \sigma_s)$. We draw a large number of managers from this distribution, have those with sufficiently high reputations raise capital and start funds with optimal lockups, and simulate five years of data. We record the distribution of optimal lockup maturities and the survival rates. We repeat these simulations to find μ_s and σ_s by matching, between the simulations and the data, this distribution of lockups and the survival rates.

The model matches these features of the data when the managers' log-likelihoods of reputations are drawn from a distribution with a mean of $\mu_s = 0.0625$ and a standard deviation of $\sigma_s = 2.08$. With this distribution of reputation, the first-year attrition rate is 17.6%, and the fractions of managers choosing contracts with 0- to 3-month, 3- to 12-month, and longer-than-12-month maturities are 23%, 60%, and 17%. In the data, the first-year attrition rate is 18% (Brown, Goetzmann, and Park 2001) and the

fractions of funds with these lockup maturities are 30%, 67%, and 3% (Hedge Fund Research (HFR) database).⁸

In the model investors have rational expectations, and so the distribution of reputations also tells the fraction of skilled managers. With this distribution of reputation, 16% of potential managers have sufficiently high reputations to open a fund. Some skilled managers are below the threshold, and some unskilled managers are above it. A total of 40% of all skilled managers have too low reputations to raise capital, and the share of skilled managers among those who start a fund is 58%. This proportion increases over time as investors learn about returns and liquidate managers they come to perceive as unskilled. In a year, the proportion of skilled managers among funds that remain alive has increased to 66%; investors rationally update their beliefs and therefore predominantly liquidate unskilled managers. The proportion of skilled managers reaches 80% after three years.

Because many skilled managers are only marginal in terms of their reputations, they initially distort their behavior significantly. In the beginning, the average skilled manager forgoes 78 basis points in expected long-term returns to signal skill by enhancing short-term returns. This is the unconditional estimate that averages across the crash states. If the economy starts in the low crash state, there is no distortion; and if it starts in the high crash state, the distortion is exactly twice as high at 156 basis points. These distortions dissipate over time as the reputations of managers near the liquidation threshold either fall down (and the managers get liquidated) or move up. Among managers who survive a year, the average distortion falls from 78 to 29 basis points. These distortions occur with managers offering their investors contracts with optimal lockup maturities. If skilled managers could only offer open-ended contracts, the initial distortions would be higher at 167 and 332 basis points, respectively.

⁸Three different hedge fund contract terms limit investors ability to withdraw funds and are therefore equivalent to lockups in our model: lockups, advance notices, and redemption wait times. In classifying funds in the HFR database, we set each fund's lockup maturity equal to the maximum of these values.

Lockups therefore substantially reduce, but do not fully eliminate, the distortions that emerge from the signaling problem.

Although these estimates are based on a stylized model, our approach makes a more general point. Because optimal contracts and attrition rates depend on investors' perceptions of skill, we can draw inferences about skill from a model such as ours without estimating fund alphas. This ability is particularly valuable in the hedge fund industry. Hedge funds are exempt from the Investment Company Act of 1940 and therefore not subject to its disclosure requirements. As a consequence, hedge fund databases are subject to various biases, such as the management of reported returns, backfill and incubation biases, survivorship bias, liquidation bias, and self-reporting bias.⁹ Using a structural model to estimate the distribution of skill is complementary to estimating alphas from these hedge fund databases.

6 Conclusions

We present a dynamic model in which rational investor learning leads managers to improve their short-term returns at the expense of expected long-term returns. The model highlights a fundamental delegation friction that stems from investors' inability to distinguish, in the short-run, a profitable long-term strategy from lack of skill. We show that this friction generates large limits to arbitrage. Managers do not trade the arbitrage opportunity to the full extent because they are concerned about investors interpreting their low short-run returns as evidence of lack of skill. Limits to arbitrage is a robust feature

⁹See, for example, Bollen and Pool (2009) and Agarwal, Daniel, and Naik (2011) on the management of reported returns; Malkiel and Saha (2005) on backfill and incubation biases; Fung and Hsieh (1997, 2000) and Brown, Goetzmann, and Park (2001) for survivorship bias; Ackermann, McEnally, and Ravenscraft (1999) for liquidation bias; and Agarwal, Fos, and Jiang (2013) for self-reporting bias. Backfill bias refers to hedge funds ability to begin reporting returns only if their performance has been good, and to backfill returns that they earned before entering the platform. Incubation bias refers to the fact that it mainly the successful funds that show up on the platforms; see, also, Evans (2010). Survivorship bias refers to the bias in returns that occurs if the databases do not contain information on all liquidated funds. Liquidation bias refers to the bias that emerges when funds stop reporting their (low) returns when they are about to get liquidated. Self-reporting bias refers to hedge funds' ability to choose not to report to any platform or to discontinue and later continue reporting for various reasons.

of an environment in which investors are uncertain about managerial skill, and managers have to decide how aggressively to trade the long-term arbitrage opportunity. A wide range of parametrizations leads to quantitatively large limits to arbitrage.

Lockups, with similar maturities to those in the hedge fund data, reduce, but do not eliminate, limits to arbitrage. Managers benefit from the lockups. With investors committed not to liquidate, managers can trade the long-term arbitrage opportunity more aggressively without worrying about the signal sent by low short-term returns. Investors, on the other hand, perceive these lockups as costly. They are ex ante concerned about the possibility of being locked up with an unskilled manager who destroys value. The optimal lockup maturity strikes a balance between these tradeoffs.

The model tells us to ask four questions when evaluating or designing a management contract. What is the typical horizon of the manager investment? How informative is short-term performance about the managerial skill? How competitive is the market for skill? What is the cost of lack of skill? The answers to these questions determine the extent to which a lockup provision can benefit both managers and investors.

Although we discuss the model in terms of, and calibrate it to, hedge funds, its fundamental friction applies to other contracts, such as CEO compensation. Whenever there is uncertainty about managerial skill and when profitable long-term opportunities are sometimes costly in the short run, the insights of our model apply.

A Appendix

B Propositions

Proposition 1. *Evolution of beliefs.* The log-likelihood ratio of investors' beliefs about skill, $p_t = \log\left(\frac{P_t}{1-P_t}\right)$, evolves as

$$dp_t = \frac{\kappa_t}{\sigma} (dR_t - \check{\mu}_t dt) + \frac{\kappa_t^c(\xi_t)}{\omega} (dR_{t+} - \check{\mu}_t^c(\xi_t)). \quad (13)$$

In equation (13), dR_t is the instantaneous non-crash return, dR_{t+} is the crash-event return, κ_t and $\kappa_t^c(\xi_t)$ are the signal-to-noise ratios of short-term performance and crash performance, and $\check{\mu}_t$ and $\check{\mu}_t^c(\xi_t)$ are investors' expectations about short-term and crash performance in log-likelihood space:

$$\begin{aligned} \kappa_t &= \frac{(\alpha^g + \frac{\hat{\pi}_t(\xi^h) + \hat{\pi}_t(\xi^l)}{2} \lambda) - (\alpha^b + \lambda)}{\sigma}, \\ \kappa_t^c(\xi) &= \frac{-\xi(\hat{\pi}_t(\xi) - 1)}{\omega}, \\ \check{\mu}_t &= \frac{(\alpha^g + \frac{\hat{\pi}_t(\xi^h) + \hat{\pi}_t(\xi^l)}{2} \lambda) + (\alpha^b + \lambda)}{2}, \text{ and} \\ \check{\mu}_t^c(\xi) &= -\xi \frac{\hat{\pi}_t(\xi) + 1}{2}. \end{aligned}$$

Proof. Let μ_t^g and μ_t^b denote investors' expectations about skilled and unskilled managers' short-term performance at time t when there is no crash. Also let $R_{t,\Delta} = \int_t^{t+\Delta} dR_t$. The problem facing the analysts is that of distinguishing between two normal distributions with different means but equal variances. Bayes rule implies that if there is no crash between time t and $t + \Delta$, investors' posterior belief at time $t + \Delta$ equals

$$P_{t+\Delta} = \frac{e^{-\frac{(R_{t,\Delta} - \mu_t^g \Delta)^2}{2\sigma^2 \Delta}} P_t}{e^{-\frac{(R_{t,\Delta} - \mu_t^g \Delta)^2}{2\sigma^2 \Delta}} P_t + e^{-\frac{(R_{t,\Delta} - \mu_t^b \Delta)^2}{2\sigma^2 \Delta}} (1 - P_t)}. \quad (14)$$

The log-likelihood ratio $p_t = \frac{P_t}{1-P_t}$ therefore evolves as

$$P_{t+\Delta} = \log \left(\frac{e^{-\frac{(R_{t,\Delta} - \mu_t^g \Delta)^2}{2\sigma^2 \Delta}} P_t}{e^{-\frac{(R_{t,\Delta} - \mu_t^b \Delta)^2}{2\sigma^2 \Delta}} (1 - P_t)} \right) = P_t + \frac{\mu_t^g - \mu_t^b}{\sigma} \times \frac{1}{\sigma \Delta} \left(R_{t,\Delta} - \frac{\mu_t^g + \mu_t^b}{2} \Delta \right). \quad (15)$$

Investors' expectations about short-term performance depend on their beliefs about the skilled manager's portfolio choice, $\hat{\pi}_t$. Let κ_t denote the signal-to-noise ratio of short-term performance and $\check{\mu}_t$ the average short-term performance across the manager types (the investor expectation in log-likelihood space):

$$\begin{aligned} \kappa_t &= \frac{\mathbf{E}_t(dR_t|\theta = g) - \mathbf{E}_t(dR_t|\theta = b)}{\sigma} = \frac{(\alpha^g + \frac{\hat{\pi}_t(\xi^h) + \hat{\pi}_t(\xi^l)}{2} \lambda) - (\alpha^b + \lambda)}{\sigma}, \text{ and} \\ \check{\mu}_t &= \frac{1}{2}(\mathbf{E}_t(dR_t|\theta = g) + \mathbf{E}_t(dR_t|\theta = b)) = \frac{1}{2} \left((\alpha^g + \frac{\hat{\pi}_t(\xi^h) + \hat{\pi}_t(\xi^l)}{2} \lambda) + (\alpha^b + \lambda) \right). \end{aligned}$$

The log-likelihood process of perceived skill, in absence of crashes, therefore evolves as given by the first term of Equation (13).

When there is a crash, the manager experiences a large return. Because the unskilled manager is always long the arbitrage opportunity, he loses $-\xi_t$ when the crash hits. The skilled manager's return depends on $\hat{\pi}_t(\xi_t)$. Investors again need to distinguish between two normal distributions that have different means (if $\hat{\pi}_t \neq 1$) but equal variances; the return volatility associated with the crash is ω . Let $\kappa_t^c(\xi)$ denote the signal-to-noise ratio of crash performance in a given crash state and $\check{\mu}_t^c(\xi)$ the average crash performance across the manager types in a given crash state:

$$\begin{aligned} \kappa_t^c(\xi) &= \frac{\mathbf{E}_t(dR_{t+}|\theta = g, \xi_t = \xi) - \mathbf{E}_t(dR_{t+}|\theta = b, \xi_t = \xi)}{\omega} = \frac{-\xi(\hat{\pi}_t(\xi) - 1)}{\omega}, \text{ and} \\ \check{\mu}_t^c(\xi) &= \frac{1}{2}(\mathbf{E}_t(dR_{t+}|\theta = g, \xi_t = \xi) + \mathbf{E}_t(dR_{t+}|\theta = b, \xi_t = \xi)) = \frac{1}{2}(-\xi(\hat{\pi}_t(\xi) + 1)). \end{aligned}$$

Applying Equation (15) to the crash return, we obtain the second term of Equation (13). \square

Proposition 2. A manager's valuation increases in his reputation. Let $\delta_c = 0$, $\underline{G} = 0$, $f > 0$, $\kappa_t = \kappa > 0$, and $\kappa_t^c(\xi) = \kappa^c(\xi) > 0$, and fix the liquidation threshold $p_L(L)$ for the two lockup states, as well as the skilled manager's portfolio policy $\pi_t(\xi) = \pi(\xi)$. The manager's valuation then increases in his reputation, $G_p > 0$.

Proof. If no outside offers are made ($\delta_c = 0$), then the manager reputation only impacts her valuation through the liquidation event. The conditions on fees and manager skill imply that the manager has more value alive than when liquidated. Holding the manager's portfolio and investors' beliefs constant, then for reputations $p_1 > p_2$ and over an interval Δ , $\Pr(p_{t+\Delta} < p_L | p_1) < \Pr(p_{t+\Delta} < p_L | p_2)$.

The distribution of $p_{t+\Delta}$ conditional on p_1 first-order stochastically dominates the distribution conditional on p_2 . The probability of fund liquidation is given by the first time the reputation reaches threshold p_L when the lockup has expired. Since the lockup expiration process is independent of the manager's initial reputation, the probability of fund liquidation is increasing in probability of hitting p_L . Since $p_{t+\Delta} | p_1$ FOSD $p_{t+\Delta} | p_2$, it follows that the hitting time probability is weakly smaller for p_1 than for p_2 . It is intuitive that $p_{t+\Delta} | p_1$ FOSD $p_{t+\Delta} | p_2$, since both share the same distribution. Formally, $p_{t+\Delta} | p$ can be written as a mixture of normally distributed random variables (Ait-Sahalia 2004):

$$\begin{aligned} \text{Prob}(p_{t+\Delta} < z | p, \xi) &= \sum_{n=0}^1 \frac{e^{-\delta_\xi \Delta} (\delta_\xi \Delta)^n}{n!} \Phi \left(\frac{z - A(n, n, 0, \Delta, \xi) - p}{B(n, n, 0, \Delta, \xi)} \right) \\ &+ \sum_{n=2}^{\infty} \sum_{n_h=0}^{n-1} \binom{n-1}{n_h} \left(\frac{1}{2} \right)^{n_h} \left(1 - \frac{1}{2} \right)^{n-n_h} \frac{e^{-\delta_\xi \Delta} (\delta_\xi \Delta)^n}{n!} \times \\ &\Phi \left(\frac{z - \mu(1, n, n_h, \Delta, \xi) - p}{B(1, n, n_h, \Delta, \xi)} \right), \end{aligned}$$

where $\bar{\pi} = \frac{1}{2}(\pi(\xi^h) + \pi(\xi^l))$ and

$$A(n_0, n, n_h, \Delta, \xi) = \left(\frac{\kappa}{\sigma} (\alpha^s + \bar{\pi}\lambda - \check{\mu}) \Delta - n_0 \frac{\kappa^c(\xi)}{\omega} (-\pi(\xi)\xi - \check{\mu}^c(\xi)) + n_h \frac{\kappa^c(\xi^h)}{\omega} (-\pi(\xi^h)\xi^h - \check{\mu}^c(\xi^h)) \right. \\ \left. + (n - n_0 - n_h) \frac{\kappa^c(\xi^l)}{\omega} (-\pi(\xi^l)\xi^l - \check{\mu}^c(\xi^l)) \right) \quad (16)$$

$$B(n_0, n, n_h, \Delta, \xi) = \sqrt{n_h (\kappa^c(\xi^h))^2 + (n - n_0 - n_h) (\kappa^c(\xi^l))^2 + n_0 (\kappa^c(\xi))^2 + \Delta \kappa^2} \quad (17)$$

Since for each of these normal distributions we have $\Phi\left(\frac{z - A(n_0, n, n_h, \Delta, \xi) - p_1}{B(n_0, n, n_h, \Delta, \xi)}\right) < \Phi\left(\frac{z - A(n_0, n, n_h, \Delta, \xi) - p_2}{B(n_0, n, n_h, \Delta, \xi)}\right)$ if $p_1 > p_2$, it follows that $Prob(p_{t+\Delta} < z | p_1, \xi) < Prob(p_{t+\Delta} < z | p_2, \xi)$. Under these assumption we can rewrite the manager valuation in (7) as $G = \int_0^\infty e^{-rt} E[\gamma(dR_t - (f + r)dt) + f dt | p] (1 - \Pr(\tau_l < t | p)) dt$, that is the sum of the discounted expected cash-flows weighted by the probability of being alive. Since the manager's portfolio choice is assumed to be fixed across reputations, the expected cash flows are independent of p . The manager's valuation therefore depends (negatively) on the probability of fund liquidation. Since we proved that this probability is decreasing in reputation, it follows that manager valuations are increasing in reputations ($G_p > 0$). \square

Proposition 3. *A manager's valuation is S-shaped in his reputation. Let $\delta_c = 0$, $\underline{G} = 0$, $f > 0$, $\pi_t(\xi) = \pi$, and $\omega \rightarrow \infty$, and fix the liquidation threshold $p_L(L)$ for the two lockup and crash states as well as the manager's portfolio. Then,*

(1) *A skilled manager who manages an open-ended fund is risk-averse with respect to reputation risk:*

$$p_t \geq p_L \Rightarrow G_{pp} < 0.$$

(2) *A skilled manager who manages a fund with an active lockup is risk-averse with respect to reputation risk when his reputation is sufficiently high, and risk-loving when his reputation is low:*

$$p_t \geq p^*(T) \Rightarrow G_{pp} < 0 \text{ and } p_t < p^*(T) \Rightarrow G_{pp} > 0, \text{ where } p^*(T) \text{ is a critical threshold that}$$

depends on the length of the lockup provision.

Proof. When $\omega \rightarrow 0$ learning is perfect during crashes. The results hold more generally, but the “no learning in crashes” case is convenient as the manager HJB becomes a system of four ODE’s with constant coefficients, with boundary condition $G(p_t, L_t = 0, \xi_t) = 0$ for $p \leq p_L(L = 0)$. The constructive proof is almost identical in the opposite extreme, i.e. if learning is perfect, that is if the manager reputation jumps to $p = \infty$ the first time a crash arrives. The four different ODEs are for states $\{\xi_t, L_t\} \in (\{\xi^h, \xi^l\} \times \{0, 1\})$,

$$\begin{aligned} 0 = & \gamma(\alpha^g + \pi(\lambda - \delta_\xi \xi_t) - f) + f - rG(t) + G_p(t) \frac{\kappa_t}{\sigma} (\alpha^g + \pi\lambda - \check{\mu}_t) + \frac{1}{2} G_{pp}(t) (\kappa_t)^2 \\ & + \delta_e (G(X_t, L_t = 0, \xi_t) - G(t)) + \delta_\xi \left(E^\xi [G(p_t, L_t, \tilde{\xi})] - G \right) \end{aligned}$$

Under the assumptions of the proposition, the only difference between the ODEs are the ξ state that shows up only in the first term and the transition from a lockup fund to an open-ended fund. We start with the open-ended fund. Since there are no outside offers, an open-ended fund never transitions back to a lockup fund, and so the ODE system can be decoupled from the lockup ODEs:

$$\begin{aligned} 0 = & \gamma(\alpha^g + \pi(\lambda - \delta_\xi \xi^h) - f) + f - rG(p, \xi^h) + G_p(p, \xi^h) \frac{\kappa_t}{\sigma} (\alpha^g + \pi\lambda - \check{\mu}_t) + \frac{1}{2} G_{pp}(p, \xi^h) (\kappa_t)^2 + \frac{\delta_\xi}{2} \left(G(p, \xi^l) - G(p, \xi^h) \right) \\ 0 = & \gamma(\alpha^g + \pi(\lambda - \delta_\xi \xi^l) - f) + f - rG(p, \xi^l) + G_p(p, \xi^l) \frac{\kappa_t}{\sigma} (\alpha^g + \pi\lambda - \check{\mu}_t) + \frac{1}{2} G_{pp}(p, \xi^l) (\kappa_t)^2 + \frac{\delta_\xi}{2} \left(G(p, \xi^h) - G(p, \xi^l) \right) \end{aligned}$$

Let $W(p) = G(p, \xi^l) - G(p, \xi^h)$. Then, subtracting both equations we obtain

$$0 = \gamma \delta_\xi \pi (\xi^h - \xi^l) - rW(p) + W_p(p) \frac{\kappa}{\sigma} (\alpha^g + \pi\lambda - \check{\mu}) + \frac{1}{2} W_{pp}(p) (\kappa)^2 - \delta_\xi W(p)$$

Check that $W(p) = \frac{\gamma \pi \delta_\xi (\xi^h - \xi^l)}{r + \delta_\xi}$, satisfies this equation. This tells us that $G_p(p, \xi^l) = \frac{\gamma \pi \delta_\xi (\xi^h - \xi^l)}{r + \delta_\xi} + G_p(p, \xi^h)$. The system simplifies to one ODE,

$$0 = \gamma(\alpha^g + \pi(\lambda - \delta_\xi \xi^h) - f) + f - rG(p, \xi^h) + G_p(p, \xi^h) \frac{\kappa}{\sigma} (\alpha^g + \pi\lambda - \tilde{\mu}) + \frac{1}{2} G_{pp}(p, \xi^h) (\kappa)^2 + \delta_\xi 0.5 \left(\frac{\gamma\pi\delta_\xi(\xi^h - \xi^l)}{r + \delta_\xi} \right)$$

Which homogenous solution can be easily found to be:

$$G(p, \xi^h) = \frac{\gamma(\alpha^g + \pi(\lambda - \delta_\xi \xi^h) - f) + f + \frac{\delta_\xi}{2} \frac{\gamma\pi\delta_\xi(\xi^h - \xi^l)}{r + \delta_\xi}}{r} + K_1 e^{X\eta_1} + K_2 e^{x\eta_2}$$

with $\eta_1 < 0 < \eta_2$. There are two relevant boundary conditions $G(p_L, \xi^h) = 0$ and

$$\lim_{p \rightarrow \infty} G(p, \xi^h) = \frac{\gamma(\alpha^g + \pi(\lambda - \delta_\xi \xi^h) - f) + f + 0.5\delta_\xi \frac{\gamma\pi\delta_\xi(\xi^h - \xi^l)}{r + \delta_\xi}}{r}, \quad (18)$$

the value that the manager would earn if his reputation is perfect and therefore he is never liquidated.

The second boundary implies $K_2 = 0$, since $\eta_2 > 0$. K_1 is determined by:

$$\frac{\gamma(\alpha^g + \pi(\lambda - \delta_\xi \xi^h) - f) + f + 0.5\delta_\xi \frac{k\pi\delta_\xi(\xi^h - \xi^l)}{r + \delta_\xi}}{r} + K_1 e^{p_L \eta_1} = 0$$

then

$$G(p, \xi^h, L_t = 0) = \begin{cases} \frac{\gamma(\alpha^g + \pi(\lambda - \delta_\xi \xi^h) - f) + f + 0.5\delta_\xi \frac{\gamma\pi\delta_\xi(\xi^h - \xi^l)}{r + \delta_\xi}}{r} (1 - e^{(p - p_L)\eta_1}) & p > p_L \\ 0 & X \leq p_L \end{cases}$$

It is trivial to check that $\text{Sign}[G_{pp}] = \text{Sign}[G \times (-n_1^2)] < 0$. So managers of open-ended funds are always averse to reputational risk. For funds with lockups, we have

$$\begin{aligned}
0 &= \gamma(\alpha^g + \pi(\lambda - \delta_\xi \xi^h) - f) + f - rG + G_p \frac{\kappa}{\sigma} (\alpha^g + \pi\lambda - \check{\mu}) + \frac{1}{2} G_{pp} (\kappa_t)^2 \\
&\quad + \delta_\xi 0.5 (G(p, \xi^l, L_t = 1) - G(p, \xi^h, L_t = 1)) + \delta_e (G(p, \xi^h, L_t = 0) - G(p, \xi^h, L_t = 1)) \\
0 &= \gamma(\alpha^g + \pi(\lambda - \delta_\xi \xi^l) - f) + f - rG + G_p \frac{\kappa}{\sigma} (\alpha^g + \pi\lambda - \check{\mu}) + \frac{1}{2} G_{pp} (\kappa)^2 \\
&\quad + \delta_\xi 0.5 (G(p, \xi^h, L_t = 1) - G(p, \xi^l, L_t = 1)) + \delta_e (G(p, \xi^l, L_t = 0) - G(p, \xi^l, L_t = 1))
\end{aligned}$$

Following identical steps as for open-ended funds, we have that $G(p, \xi^l, L_t = 1) - G(p, \xi^h, L_t = 1) = \frac{\kappa \pi \delta_\xi (\xi^h - \xi^l)}{r + \delta_\xi}$, plugging in the equation above we obtain

$$\begin{aligned}
0 &= \gamma(\alpha^g + \pi(\lambda - \delta_\xi \xi^h) - f) + f - rG + G_p \frac{\kappa}{\sigma} (\alpha^g + \pi\lambda - \check{\mu}) + \frac{1}{2} G_{pp} (\kappa)^2 \\
&\quad + \delta_\xi 0.5 \left(\frac{\gamma \pi \delta_\xi (\xi^h - \xi^l)}{r + \delta_\xi} \right) + \delta_e (G(p, \xi^h, L_t = 0) - G(p, \xi^h, L_t = 1)).
\end{aligned}$$

Below the liquidation threshold, $p \leq p_L(0)$ we have $G(p, \xi^h, L_t = 0) = 0$, then

$$G(p, \xi^h, L = 1) = \frac{\gamma(\alpha^g + \pi(\lambda - \delta_\xi \xi^h) - f) + f + 0.5 \delta_\xi \frac{\gamma \pi \delta_\xi (\xi^h - \xi^l)}{r + \delta_\xi}}{r + \delta_e} + K_1 e^{X \eta_3} + K_2 e^{x \eta_4}$$

with boundary conditions $\lim_{p \rightarrow -\infty} G(p, \xi^h, L = 1) = \frac{\gamma(\alpha^g + \pi(\lambda - \delta_\xi \xi^h) - f) + f + 0.5 \delta_\xi \frac{\gamma \pi \delta_\xi (\xi^h - \xi^l)}{r + \delta_\xi}}{r + \delta_e}$, the value of managing the fund if the manager is liquidated for sure once the lockup expires, this implies $K_1 = 0$. So below p_L we have that $\text{Sign}[G_{pp}] = \text{Sign}[K_2 \times (n_4^2)]$. Because $\eta_4 > 0$ and $K_2 > 0$ and since we know that $G(p, \xi^h, L_t = 1)$ is increasing in reputation, we have that $\text{Sign}[G_{XX}] > 0$. So managers with reputations below the liquidation threshold like reputational risk when they have active lockups. We use the same strategy to solve for the fund lockup value above the threshold and solve the valuation of the difference

between the value function. Let $W(p, \xi^h) = G(p, \xi^h, L_t = 1) - G(p, \xi^h, L_t = 0)$

$$0 = -rW(p, \xi^h) + W_p \frac{\kappa}{\sigma} (\alpha + \pi\lambda - \check{\mu}) + \frac{1}{2} W_{pp} (\kappa)^2 - \delta_e W(p, \xi^h)$$

which solution is of the form:

$$W(p, \xi^h) = K_3 e^{p\eta_3} + K_4 e^{p\eta_4}$$

and the roots $\eta_3 < 0 < \eta_4$ are the same roots from the value function below the threshold. So above the threshold we have $G(p, \xi^h, L_t = 1) = G(p, \xi^h, L_t = 0) + K_3 e^{X\eta_3} + K_4 e^{X\eta_4}$, since as $p \rightarrow \infty$ the manager is not liquidated, the value of the open-ended fund has to converge to the value of the lockup fund $\lim_{p \rightarrow \infty} G(p, \xi^h, L_t = 1) = G(p, \xi^h, L_t = 0)$, this implies $K_4 = 0$. Let $G_\infty = \gamma(\alpha^g + \pi(\lambda - \delta_\xi \xi^h) - f) + f + 0.5\delta_\xi \frac{\gamma\pi\delta_\xi(\xi^h - \xi^l)}{r + \delta_\xi}$ be the value of the infinite lived manager, then the value of a lockup fund is of the form

$$G(p, \xi^h, L_t = 1) = \begin{cases} \frac{G_\infty}{r} (1 - e^{(p-p_L)\eta_1}) + K_3 e^{p\eta_3} & p > p_L, \\ \frac{G_\infty}{r + \delta_e} + K_2 e^{p\eta_4} & p \leq p_L. \end{cases}$$

We can solve K_3 and K_2 by imposing value-matching ($\lim_{p \nearrow p_L} G(p, \xi^h, L_t = 1) = \lim_{p \searrow p_L} G(p, \xi^h, L_t = 1)$) and smooth-pasting ($\lim_{p \nearrow p_L} G_p(p, \xi^h, L_t = 1) = \lim_{p \searrow p_L} G_p(p, \xi^h, L_t = 1)$). See Dixit (1993) for a discussion why these are the relevant boundary conditions for a transitional boundary. This gives us a system of two equations and two unknowns which yields,

$$\begin{aligned} K_3 &= e^{-p_L \eta_3} \left(\frac{G_\infty}{r + \delta_e} + \frac{-\frac{G_\infty}{r} \eta_1 + \eta_3 \left(\frac{G_\infty}{r + \delta_e} \right)}{(\eta_4 - \eta_3)} \right) \\ K_2 &= e^{-p_L \eta_4} \frac{-\frac{G_\infty}{r} \eta_1 + \eta_3 \left(\frac{G_\infty}{r + \delta_e} \right)}{(\eta_4 - \eta_3)} \end{aligned}$$

Above p_L , $G_{pp} = -\frac{G_\infty}{r}(\eta_1^2 e^{(p-p_L)\eta_1}) + \eta_3^2 e^{(p-p_L)\eta_3} \left(\frac{G_\infty}{r+\delta_e} + \frac{-\frac{G_\infty}{r}\eta_1 + \eta_3(\frac{G_\infty}{r+\delta_e})}{(\eta_4 - \eta_3)} \right)$. Since it can be shown that $\left(\frac{G_\infty}{r+\delta_e} + \frac{-\frac{G_\infty}{r}\eta_1 + \eta_3(\frac{G_\infty}{r+\delta_e})}{(\eta_4 - \eta_3)} \right) > 0$ and $\eta_1, \eta_3 < 0$, we see that the risk aversion can alternate between positive and negative depending on which root dominates. These roots are function of the model primitives and are given by:

$$\begin{aligned} \varrho(\delta_e) &= \sqrt{a^2 + 2(r + \delta_e)\sigma^2}, \\ \eta_1 &= \frac{-a - \varrho(0)}{\kappa_t \sigma}, \\ \eta_3 &= \frac{-a - \varrho(\delta_e)}{\kappa_t \sigma}, \\ \eta_4 &= \frac{-a + \varrho(\delta_e)}{\kappa_t \sigma}, \end{aligned}$$

where $a = (\alpha^g + \pi\lambda - \check{\mu})$. Because $\eta_3 < \eta_1 < 0 < \eta_4$, the risk-aversion term will dominate when reputation is sufficiently high: $\lim_{p \rightarrow \infty} \frac{G_{pp}(p, \xi^h, L_t=1)}{e^{(p-p_L)\eta_1}} = -\frac{G_\infty}{r}\eta_1^2 \leq 0$. Since $\eta_3 < \eta_1 < 0$, either $|\frac{G_\infty}{r}(\eta_1^2)| \geq \eta_3^2 \left(\frac{G_\infty}{r+\delta_e} + \frac{-\frac{G_\infty}{r}\eta_1 + \eta_3(\frac{G_\infty}{r+\delta_e})}{(\eta_4 - \eta_3)} \right)$, in which case the manager is risk averse for any $p > p_L$, or $|\frac{G_\infty}{r}(\eta_1^2)| < \eta_3^2 \left(\frac{G_\infty}{r+\delta_e} + \frac{-\frac{G_\infty}{r}\eta_1 + \eta_3(\frac{G_\infty}{r+\delta_e})}{(\eta_4 - \eta_3)} \right)$, in which case the manager is risk-loving for values above p_L and below p^* . The threshold p^* is defined by:

$$p^* = p_L + \frac{1}{\eta_1 - \eta_3} \ln \left(\frac{r}{G_\infty} \frac{\eta_3^2}{\eta_1^2} \left(\frac{G_\infty}{r + \delta_e} + \frac{-\frac{G_\infty}{r}\eta_1 + \eta_3(\frac{G_\infty}{r+\delta_e})}{(\eta_4 - \eta_3)} \right) \right).$$

□

Proposition 4. *Longer lockup provisions reduce short-term reputational concerns.* Let $T = 1/\delta_e$ be the length of the lockup provision. Under the same conditions as Proposition 3, and if idiosyncratic volatility σ is not too large, then for any $p > p_L$, $\frac{\partial G_p(p, \xi|T)}{\partial T} > 0$.

Proof. From Proposition 3, we have

$$G_p = -\eta_1 \frac{G_\infty}{r} e^{(p-p_L)\eta_1} + \eta_3(\delta_e) \left(\frac{G_\infty}{r + \delta_e} + \frac{-\frac{G_\infty}{r} + \eta_3(\delta_e) \frac{G_\infty}{r + \delta_e}}{\eta_4(\delta_e) - \eta_3(\delta_e)} \right) e^{(p-p_L)\eta_3}, \quad (19)$$

where we have made explicit that roots η_3 and η_4 depend on the length of the lockup. We need to show that $\frac{\partial G_p}{\partial \delta_e} > 0$ to prove the result. Differentiating the above expression,

$$\begin{aligned} \frac{\partial G_p(p, \xi, L = 1)}{\partial \delta_e} &= (p - p_L) e^{-\frac{(p-p_L)(a+e[\delta_e])}{\kappa\sigma}} G(p, \xi, L = 1) (a^3 + 2(\delta_e + r)a\sigma^2 + (a^2 + 2(\delta_e + r)\sigma^2)\varrho[0]) \\ &\quad + \frac{1 - e^{-\frac{(p-p_L)(a+e[\delta_e])}{\kappa\sigma}} G(p, \xi, L = 1) \kappa_t \sigma (a(a + \varrho[0]) - 2r\sigma^2)}{2r\kappa^2(a^2 + 2(\delta_e + r)\sigma^2)^{3/2}}. \end{aligned}$$

Note that the first term dominates when reputation is high, while for reputation close to the liquidation threshold the second term dominates. The first term is positive as long as $a > 0$. So $\frac{\partial G_p}{\partial \delta_e} > 0 \Rightarrow \frac{\partial G_p}{\partial T} < 0$ for high enough reputations. The second term has a positive constant term and a negative term that is decreasing (in absolute value) in the reputation. Note that this second term goes to zero as the lockup maturity goes to zero ($\delta_e \rightarrow \infty$). Reputational concerns therefore decrease with increase in lockups when lockups are short-enough, there is a $T^* > 0$, such that for any $T < T^*$, $\frac{\partial G_p(\cdot|T)}{\partial T} < 0$. So we have shown so far that reputation concerns are decreasing in lockup maturity for any maturity for high enough reputations, and for any reputation for short enough maturities. To prove the general result, we need conditions under which

$$G_\infty \kappa \sigma^3 \left(a \frac{a + \sqrt{a^2 + 2r\sigma^2}}{\sigma^2} - 2r \right) < 1.$$

This condition holds for reasonable parameters. For example, with volatility of 0.1, signal-to-noise ratio of 1, drift of 0.1, and a total fund value V of 10, this expression has value ≈ 0.01 . We can also frame this as a condition on the amount of idiosyncratic risk. Recall that $\kappa = \frac{\alpha^g + \pi\lambda - (\lambda + \alpha^b)}{\sigma}$. The

expression is clearly increasing in σ and holds for $\sigma \rightarrow 0$. So given the model primitives we can define σ^* implicitly as

$$G_\infty(\alpha^g + \pi\lambda - (\lambda + \alpha^b)) \left(a(a + \sqrt{a^2 + 2r\sigma^{*2}}) - \frac{2r}{\sigma^{*2}} \right) = 1$$

So $\frac{\partial G_p(\cdot|T)}{\partial T} < 0$ for any $\sigma < \sigma^*$. This condition is not an important limitation because everything else constant reputation concerns are *decreasing* in idiosyncratic volatility. So lockups reduce reputation concerns when reputation concerns are large. \square

Proposition 5. Manager's equilibrium portfolio choice. *Let $\kappa > 0$, $G \geq 0$, and $G_p \geq 0$. We let $\pi(p_t, L_t, \xi_t)$ denote the skilled manager's optimal portfolio choice. In the equilibrium defined in Section 2.5,*

$$\pi(p_t, L_t, \xi_t) = \begin{cases} 1 & \text{if } \xi_t = \xi^l, \\ g(p_t, L_t) & \text{if } \xi_t = \xi^h, \end{cases} \quad (20)$$

where $g(p_t, L_t)$ solves:

$$g = \max_g g(\lambda - \delta_\xi \xi^h), \quad (21)$$

$$\begin{aligned} \text{s.t. } g = \arg \max_{\pi \in [-1, 1]} & \left\{ \gamma \pi(\lambda - \delta_\xi \xi^h) + G_p \frac{\alpha^g + \frac{1}{2}(g-1)\lambda - \alpha^b}{\sigma^2} \pi \lambda \right. \\ & \left. + \delta_\xi E \left[G \left(p - \frac{(g-1)\xi^h}{\omega^2} \left(\tilde{\omega} - \pi \xi^h + \frac{(g+1)}{2} \xi^h \right), L, \tilde{\xi} \right) \right] \right\}, \end{aligned} \quad (22)$$

where $E(\cdot)$ takes the expectation over the unknown state $\tilde{\xi}$ and the idiosyncratic shock $\tilde{\omega}$.

Proof. Given beliefs g , the right hand side of Equation (22) is the solution of the manager problem. Therefore Equation is a fixed point problem. In equilibrium, investors beliefs about the skilled manager portfolio in the high crash state have to be equal to the actual strategy implemented by the skilled

manager in this state. Because of strategic complementarities this fixed point problem might have multiple solutions. Intuitively, looking at Equation (22) we see that a higher g increases the incentives for the manager to choose a higher pi , so there is potential for multiplicity, specially when reputation concerns are high G_p relative to compensation incentives $\gamma(\lambda - \delta_\xi \xi^h)$.

We take a mechanism design approach for equilibrium selection and state the portfolio allocation problem as the maximization of expected returns (the surplus of the relationship) subject to being consistent with the manager investment incentives. Therefore, one can think of Equation (22) as a (local) Incentive Compatibility constraint. This approach here is natural because managers and investors can communicate, thus is natural to expect coordination in the high surplus equilibrium.

The optimal choice can be broken into three regimes:

(A) $g = -1$ if

$$\left| \gamma(\lambda - \xi^h \delta_\xi) - 2\left(\frac{\xi^h}{\omega}\right)^2 \delta_\xi E \left[G_p \left(p + \frac{2\xi^h}{\omega^2} (\xi^h + \tilde{\omega}), L, \xi \right) \right] \right| > G_p \frac{\alpha^g - \lambda - \alpha^b}{\sigma^2} \lambda,$$

(B) if (A) does not hold, g is given by the minimum π that satisfies

$$\pi = \min \left\{ 1, 1 - \frac{\gamma(\lambda - \xi^h \delta_\xi) - G_p \frac{\alpha^g - \alpha^b}{\sigma^2} \lambda}{G_p \frac{\lambda^2}{2\sigma^2} + \frac{(\xi^h)^2}{\omega^2} \delta_\xi E \left[G_p \left(p - \frac{(\pi-1)\xi^h}{\omega^2} \left(-\frac{(\pi-1)}{2} \xi^h + \tilde{\omega} \right), L, \tilde{\xi} \right) \right]} \right\}, \quad (23)$$

This proposition follows directly from three conditions that hold for the equilibrium: (1) the manager's first-order condition, (2) investors' beliefs about the manager's portfolio choice, and (3) return maximization. In the state ξ^l , the manager's first-order condition is satisfied at $\pi(p_t, L_t, \xi_t = \xi^l) = 1$ —which maximizes expected returns—as long investors believe that the manager chooses $\pi_t = 1$. Therefore $\pi(p, L, \xi^l) = 1$ for any p and L .

The optimal choice in state $\xi_t = \xi^h$ divides into three regions depending on the intensity of repu-

tational concerns. We first check if reputation concerns are weak enough so that $\pi(p_t, L_t, \xi^h) = -1$ —the maximum expected return portfolio—satisfies the manager’s first-order condition. If it does, this choice is the equilibrium portfolio choice. If not, we find the lowest of π_t —because expected returns are decreasing in π_t —that satisfies the manager’s first-order condition. The resulting portfolio may lie in the interior— $-1 < \pi_t < 1$, or it may be in the corner, $\pi_t = 1$.

Intuition. The expression in case (A) in Proposition 5 shows that, at the efficient choice, the long-term reputational concern will be close to zero both when the long-term signal-to-noise ratio is very high or very low. If learning during crashes is very weak, then the manager’s after-crash reputation will be very close to his before-crash reputation. The marginal long- and short-term reputational concerns will therefore be of similar magnitude, but the marginal impact of portfolio changes on the after-crash reputation is small because the signal-to-noise ratio is low in this weak-learning case. On the other hand, if learning is very strong, the marginal impact of portfolio changes on the manager’s long-term reputation is large, but the marginal value of reputation is low if the manager maximizes expected returns. Since learning is strong, a manager who maximizes expected returns will likely have a very after-crash reputation. As long reputational concerns are decreasing in reputation ($G_{pp} < 0$), the marginal value of reputation after the crash will be much lower than before the crash, leading to a strong temptation to deviate towards more short-term oriented strategies.

While long-term reputational concerns will typically not be strong enough to keep the manager from distorting his choices, they will be effective in keeping the manager from fully maximizing short-term performance. Short-term reputational concerns introduce complementarity between investors’ beliefs and manager actions. The more investors expected to learn from short-term returns, i.e. the higher the κ , the higher the manager’s incentives to enhance these returns. On one hand the sensitivity of reputation to performance in a crash works exactly as for short-term performance. The better the

manager is expected to do when the crash hits, the higher the signal-to-noise ratio κ^c , and therefore the higher the incentives to perform well when the crash hits. But the nonlinearity built into long-term learning also induces more subtle reputational incentives. The long-term reputational concern is impacted not only by investors' beliefs $\hat{\pi} = g$ as in the case of short-term reputation concerns, but also by the actual manager's choice π . Inspecting the distribution of crash reputation growth in the second line of Equation (22), we can see that as the portfolio increases towards 1, the lower the expected reputation growth. As a result, the expectation puts more weight on low reputation states, i.e. states with higher reputation concerns $G_p \uparrow$. Therefore, reputation incentives depend on beliefs but also on the actual portfolio chosen by the manager. \square

Proposition 6. *Limits to arbitrage, distance to liquidation, and lockup maturity.*

Under the conditions of Propositions 4 and 3,

- (A) A skilled manager's incentive to invest in the long-term arbitrage opportunity increases in the distance to liquidation when the fund does not have a lockup provision.*
- (B) A skilled manager's incentive to invest in the long-term arbitrage opportunity increases in the distance to liquidation when the fund has a lockup provision and the manager's reputation is sufficiently high.*
- (B) A skilled manager's incentive to invest in the long-term arbitrage opportunity increases in the lockup maturity.*

Proof. Proposition 4 shows the conditions under which reputational concerns are decreasing in lockup maturity. Proposition 5 shows that when the long-term reversal strategy is the most profitable, the

manager's first-order condition is:

$$\gamma(\lambda - \xi^h \delta_\xi) + G_p \frac{\alpha^g + \frac{1}{2}(\pi - 1)\lambda - \alpha^b}{\sigma^2} \lambda + \frac{(\pi - 1)(\xi^h)^2}{\omega^2} \delta_\xi E \left[G_p \left(p - \frac{(\pi - 1)\xi^h}{\omega^2} \left(-\frac{(\pi - 1)}{2} \xi^h + \tilde{\omega} \right), L, \xi \right) \right] = 0.$$

If investors cannot learn from a crash-event returns, $\omega \rightarrow \infty$, this expression becomes:

$$\gamma(\lambda - \xi^h \delta_\xi) + G_p \frac{\alpha^g + \frac{1}{2}(\pi - 1)\lambda - \alpha^b}{\sigma^2} \lambda,$$

with the manager maximizing returns (i.e., choosing $\pi = -1$) if this value is negative. Note that $\frac{\alpha^g + \frac{1}{2}(\pi - 1)\lambda - \alpha^b}{\sigma^2} \lambda > 0$, $\gamma(\lambda - \xi^h \delta_\xi) < 0$. We therefore only need to show that G_p is decreasing in cases A-C. (A) Proposition 4 implies $G_{pp} < 0$ for any $p > p_L(0)$ in the case of open ended funds, so the result follows. (B) Proposition 4 implies $G_{pp} < 0$ for high enough reputations, so the (B) result follows. (C) Proposition 4 shows that $\frac{\partial G_p(\cdot|T)}{\partial T} < 0$. It follows that managers with longer lockups have stronger incentives to pursue the long-term arbitrage opportunity.

An identical argument holds when investors learn perfectly from crash-event returns, $\omega \rightarrow 0$. \square

Proposition 7. Investors' liquidation policies.

If investors do not learn during crashes, $\omega \rightarrow \infty$, then

- (a) *The liquidation threshold p_L increases in the arrival rate of the outside investors, δ_c .*
- (b) *The liquidation threshold p_L increases in the limits to arbitrage—that is, the difference between the highest available expected return and the skilled manager's expected return in the equilibrium—when the market for skill is sufficiently competitive.*

Proof. A more competitive market for skill means that the arrival rate of outside investors is higher, $\delta_c \uparrow$.

The liquidation policy satisfies $V(p_L, L_t = 0) = 1$. If outside investors bid a positive amount, current

investors experience a capital loss of $1 - V(p_t, L_t)$. Current investors never experience a capital gain since outside investors always bid zero when $V(p_t, L_t) < 1$. Everything else constant, an increase in the rate of offer arrival decreases the investors' valuation. Because the investor has to be indifferent between a share in the fund and one unit of cash at the liquidation threshold, it follows that the liquidation threshold has to increase.

Higher limits to arbitrage implu smaller expected returns $\hat{\mu}_t + \delta \hat{\mu}_t^c$; therefore, holding everything else constant, an increase in the limits to arbitrage also implies lower valuations. But higher limits to arbitrage also implies more learning from short-term performance $\kappa_t \uparrow$. Since the liquidation decision is an option on the manager skill, this effect increases investor valuations. As $\delta_c \rightarrow \infty$ this volatility effect goes to zero as outside investors bid on any positive news. The expected return effect dominates, that is, the liquidation threshold increases in the limits to arbitrage when the capital markets are sufficiently competitive. \square

C Numerical solution

We apply the finite-difference method to solve the integro-partial differential equations. To solve for optimal policies, we sequentially iterate until the value function converges. The two value functions—the skilled manager's G and the investors' V —and the three choice variables—the skilled manager's portfolio choice π_t and the investors' investment policy for funds with and without lockups—are determined jointly. Equilibrium offers $B(X_t, L_t)$ are determined once we have the investor valuation V . The state space consists of the manager's reputation in probability space ($\phi \in [0, 1]$), strategy state ($\xi \in \{\xi^h, \xi^l\}$), and the lockup status $L \in \{0, 1\}$. We solve in the probability space instead of the log-likelihood space.

We first hold constant the skilled manager's portfolio choice at the efficient choice $\pi(\xi^h, a, L) = -1$ and iterate to find what we denote the efficient solution $(G_{\text{ef}}, B_{\text{ef}}, V_{\text{ef}}, a_L^{\text{ef}}, a_{L_t=1}^{\text{ef}})$. The efficient liquidation

threshold a_L^{ef} is the lower bound to the equilibrium liquidation threshold a_L^{eq} . Starting from the efficient policies we iterate on the HJB, but this time solving for the optimal portfolio $\pi(\xi^h, a, L)$ policy at each step. Recall that $\pi(\xi^l, a, L) = 1, \forall a, L$.

The iteration procedure can be divided into steps:

1. Given choices π^{i-1} solve for $a_L^i, a_{L_t=1}^i, V^i$ such that $V^i(a_L^i, L) = 1, V^i(a_{L_t=1}^i, L_t = 1) = 1$, and V satisfy the investors' HJB equation.
2. Given $a_L^i, a_{L_t=1}^i, V^i$ and π^{i-1} solve for G^i such that the manager's valuation satisfies the HJB equation and $G^i(\xi, a_L^i, L) = 0$.
3. Given G^i solve for π^i using Proposition 5
4. If $|G^i - G^{i-1}| < \epsilon, |B^i - B^{i-1}| < \epsilon$, and $|V^i - V^{i-1}| < \epsilon$, stop. If not, repeat.

When we solve for the efficient solution—that is, the one in which the skilled manager commits to maximizing expected returns—we skip step 3 as the portfolio choice is always held at $\pi(\xi^h) = -1$.

D Lockup premium in a stylized model

Investors dislike lockups. If they invest in a fund with a lockup and the manager turns out to be unskilled, they will be stuck with the manager for the duration of the lockup. If the manager is skilled, he is likely to receive an outside offer, in which case the investors lose the manager despite the lockup. The lockup is therefore asymmetric; it more likely binds when the investors wish it did not.

The key benefit of the lockup provision is that it reduces limits to arbitrage. With the lockup in place, the manager can increase his expected return by trading the long-term arbitrage opportunity more aggressively. Investors, however, need to be compensated for the risk of being stuck with unskilled

managers. The difference between the expected net alpha of a fund with a lockup and an open-ended fund is the lockup premium. It measures the entrenchment costs perceived by investors.

We use a stylized example to illustrate the source of the lockup premium. Suppose that investors learn the true type of the manager in t years, and that they are asked to enter a contract with a T -year lockup provision, where $T > t$. We make the following assumptions. First, if the manager is revealed to be skilled at time t , outside investors discover and bid him away with probability p . Second, the riskless rate is zero, there are no performance fees, and the skilled and unskilled managers earn *after-fees* expected returns of r_g and r_b . With these assumptions, investors' expected return over T years is

$$E(r) = \begin{cases} (1-p)r_g T + pr_g t & \text{if the manager is skilled,} \\ r_b T & \text{if the manager is unskilled.} \end{cases} \quad (24)$$

If the manager is skilled, investors may have their capital returned to them after just t years; but if he is unskilled, they will earn low returns for $T > t$ years. In our notation, the manager is skilled with probability $\hat{\phi}_0$, and so investors break even if

$$E(r) \geq 0 \quad \Rightarrow \quad \hat{\phi}_0 \underbrace{\left[(1-p)r_g T + pr_g t \right]}_{\text{skilled}} + (1-\hat{\phi}_0) \underbrace{r_b T}_{\text{unskilled}} \geq 0. \quad (25)$$

Investors therefore break even when the manager's reputation is at least

$$\hat{\phi}_0 \geq \frac{-r_b T}{(1-p)r_g T + pr_g t - r_b T}. \quad (26)$$

Suppose now that we take a manager at the boundary and have him manage an open-ended fund

instead. Taking $T \rightarrow 0$ in equation (25), investors in such a fund earn a net alpha of

$$E(r) = \hat{\phi}_0 r_g + (1 - \hat{\phi}_0) r_b. \quad (27)$$

Substituting the boundary reputation from (26) into equation (27), we see that investors require a lockup premium Λ of

$$\Lambda = \frac{-r_b T}{(1 - p)r_g T + p r_g t - r_b T} (r_g - r_b) + r_b, \quad (28)$$

over the open-ended fund—that is, investors only invest in managers who they believe will deliver high expected returns. This lockup premium Λ increases in the contract maturity T , the cost of investing in the unskilled manager, $-r^b$, and the degree of competition in the market for skill, p . In the full model, the lockup premium is more complicated because the speed of learning depends on the manager’s choices and the model primitives, and because investors learn continuously from returns.

REFERENCES

- Ackermann, C., R. McEnally, and D. Ravenschaft (1999). The performance of hedge funds: Risk, return, and incentives. *Journal of Finance* 54(3), 833–874.
- Agarwal, V., N. D. Daniel, and N. Y. Naik (2009). Role of managerial incentives and discretion in hedge fund performance. *Journal of Finance* 64(5), 2221–2256.
- Agarwal, V., N. D. Daniel, and N. Y. Naik (2011). Do hedge funds manage their reported returns? *Review of Financial Studies* 24(10), 3281–3320.
- Agarwal, V., V. Fos, and W. Jiang (2013). Inferring reporting-related biases in hedge fund databases from hedge fund equity holdings. *Management Science* 59(6), 1271–1289.
- Agarwal, V. and N. Y. Naik (2004). Risks and portfolio decisions involving hedge funds. *Review of Financial Studies* 17(1), 63–98.
- Aït-Sahalia, Y. (2004). Disentangling diffusion from jumps. *Journal of Financial Economics* 74(3), 487–528.
- Aragon, G. O. (2007). Share restrictions and asset pricing: Evidence from the hedge fund industry. *Journal of Financial Economics* 83(1), 33–58.
- Barroso, P. and P. Santa-Clara (2015). Momentum has its moments. *Journal of Financial Economics* 116(1), 111–120.
- Basak, S., A. Pavlova, and A. Shapiro (2007). Optimal asset allocation and risk shifting in money management. *Review of Financial Studies* 20(5), 1583–1621.
- Berk, J. B. and R. C. Green (2004). Mutual fund flows and performance in rational markets. *Journal of Political Economy* 112(6), 1269–1295.
- Berk, J. B. and R. Stanton (2007). Managerial ability, compensation, and the closed-end fund discount. *Journal of Finance* 62(2), 529–556.
- Bollen, N. P. B. and V. K. Pool (2009). Do hedge fund managers misreport returns? Evidence from the pooled distribution. *Journal of Finance* 64(5), 2257–2288.
- Brown, S. J., W. N. Goetzmann, and J. Park (2001). Careers and survival: Competition and risk in the hedge fund and CTA industry. *Journal of Finance* 56(5), 1869–1886.

- Burnside, C., M. Eichenbaum, I. Kleshchelski, and S. Rebelo (2011). Do peso problems explain the returns to the carry trade? *Review of Financial Studies* 24(3), 853–891.
- Cherkes, M., J. Sagi, and R. Stanton (2009). A liquidity-based theory of closed-end funds. *Review of Financial Studies* 22(1), 257–297.
- Cochrane, J. H. (2011). Presidential address: Discount rates. *Journal of Finance* 66(4), 1047–1108.
- Dangl, T., Y. Wu, and J. Zechner (2006). Market discipline and internal governance in the mutual fund industry. *Review of Financial Studies* 21(5), 2307–2343.
- Diamond, D. W. (1991). Debt maturity structure and liquidity risk. *Quarterly Journal of Economics* 106(3), 709–737.
- Dixit, A. K. (1993). *The Art of Smooth Pasting*, Volume 55. Taylor & Francis.
- Evans, R. B. (2010). Mutual fund incubation. *Journal of Finance* 65(4), 1581–1611.
- Fung, W. and D. Hsieh (1997). Survivorship bias and investment style in the returns of CTAs: The information content of performance track records. *Journal of Portfolio Management* 24(1), 30–41.
- Fung, W. and D. Hsieh (2000). Performance characteristics of hedge funds and commodity funds: Natural vs. spurious biases. *Journal of Financial and Quantitative Analysis* 35(3), 291–307.
- Holmström, B. and M. Harris (1982). A theory of wage dynamics. *Review of Economic Studies* 49(3), 315–333.
- Jegadeesh, N. and S. Titman (1993). Returns to buying winners and selling losers: Implications for stock market efficiency. *Journal of Finance* 48(1), 65–91.
- Jurek, J. W. and E. Stafford (2015). The cost of capital for alternative investments. *Journal of Finance* 70(5), 2185–2226.
- Malkiel, B. G. and A. Saha (2005). Hedge funds: Risk and return. *Financial Analysts Journal* 61(6), 80–88.
- Mitchell, M. and T. Pulvino (2001). Characteristics of risk and return in risk arbitrage. *Journal of Finance* 56(6), 2135–2175.
- Moreira, A. and T. Muir (2016). Volatility-managed portfolios. *Journal of Finance* 72(4), 1611–1644.
- Shleifer, A. and R. Vishny (1997). Limits of arbitrage. *Journal of Finance* 52(1), 35–55.

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

Stein, J. C. (2005). Why are most funds open-end? Competition and the limits of arbitrage. *Quarterly Journal of Economics* 120(1), 247–272.