Hedge funds, signaling, and optimal lockups*

Juhani T. Linnainmaa Alan Moreira

May 2018

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 enter a long-term trade that increases expected returns but lowers short-term returns. Investors who rationally learn from returns may mistake a skilled manager who pursues the long-term trade for an unskilled manager. Skilled managers therefore have an incentive to avoid the long-term trade to enhance short-term returns. The tradeoff between the benefits of the long-term trade 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 146 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 Jaroslav Borovicka, Douglas Diamond, Michael Gofman, Lars Hansen, Zhiguo He, John Heaton, Christopher Hrdlicka, Ralph Koijen, Mathias Kronlund, Hong Liu (discussant), Asaf Manela, Tobias Moskowitz, Diogo Palhares, Benjamin Pugsley, Savina Rizova, Alexi Savov, Pietro Veronesi, Robert Vishny, and seminar and conference participants at University of Chicago, University of California at Los Angeles, City University of Hong Kong, Yale University, and 2018 American Finance Association meetings 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. Managers also employ other tools besides lockups—infrequent redemption windows, notification periods, discretionary gates, and back-load fees—to increase capital stickiness. In this paper 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 taking a position in a long-term trade. Security selection enhances short-term returns and the long-term trade increases expected returns at the expense of short-term returns—that is, it earns a negative carry. 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 from the long-term trade. They cannot, however, directly observe managers' positions, and so they sometimes abandon skilled, but unlucky, managers. We solve this model for the manager's optimal investments, investors' investment and liquidation policies, and lockup maturity.

Uncertainty about skill is an important feature of the model. Investors learn about managerial skill from returns, and optimally decide when to withdraw money. This threat of liquidation alters manager

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

behavior. A skilled manager knows that she might be unlucky—an aggressive bet on the long-term trade might not pay off—in which case investors infer that she is unskilled and withdraw their investments. The manager therefore has the incentive to improve her short-term returns at the expense of expected returns; they build reputation to ensure survival. Investors earn lower returns on their assets because skilled managers distort their investments to signal skill.

Learning and signaling generate an amplification loop. A manager who is concerned about being liquidated enhances her short-term returns by investing in the long-term trade less aggressively. Because investors understand that this is the optimal response by a skilled manager, they know that 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 limits arbitrage. Managers trade the long-term trade 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 she did not need to worry about signaling; but in equilibrium, with the liquidation threat, her optimal portfolio might earn an expected return of just 6%. In this example, signaling concerns limit arbitrage by 4%.

When we add lockup provisions into the model, we show that both managers *and* investors can find these provisions appealing because of the signaling problem. A lockup gives managers more time to demonstrate that they are skilled. With investors locked in, managers 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.

At the heart of the model is a tradeoff that determines the optimal lockup maturity. On the one hand, lockups create value by allowing skilled managers to enhance their expected returns by taking a more aggressive position in the long-term trade. On the other hand, investors perceive a lockup as potentially costly because it removes investors' ability to withdraw their investments if the manager turns out to be unskilled. Competition for managerial skill makes long-term contracts even less attractive by introducing a costly asymmetry. A manager who turns out to be skilled attracts more capital, diluting the stake of old investors. Investors who are offered a long-term contract today are concerned that this contract will bind only if the manager turns out to be unskilled; they therefore demand an up-front compensation—a lockup premium—for this risk of entrenchment.

We calibrate the model to match the salient features of the hedge fund data: the high attrition rate of young hedge funds and the distribution of the lockup maturities. We show that lockups alleviate but do not eliminate the limits to arbitrage created by the signaling problem. The average manager who opens a new fund forgoes 146 basis points in returns per year to build up reputation. This distortion weakens over time. Among those who survive a year, the average distortion falls to 80 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 long-term trade 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 trade 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

We build up the model in four steps for expositional clarity. We start by showing how investors optimally learn about manager skill given fund return dynamics. We then solve for the manager's optimal portfolio and show the emergence of a feedback loop between the manager's optimal portfolio and the investors' learning problem. It is this feedback loop which reduces expected fund returns. We then show how the introduction of lockups—or any other contractual feature that delays liquidation—dampens this feedback loop and thereby increases the equilibrium expected fund returns. We close the model by having investors value funds with different lockup provisions and solve for the optimal lockup maturity.

2.1 Managers' investment opportunities

A manager i is either skilled (denoted by $i = g$) or unskilled (denoted by $i = b$). A skilled manager generates abnormal returns by selecting securities and by entering a long-term trade. The manager always generates alpha by selecting securities, but she can choose her exposure to the long-term trade. The stock-selection strategy is exposed to idiosyncratic risk. The long-term trade is exposed to "crashes:" it earns a constant return except when there is a crash, in which case its value jumps. The idiosyncratic risk of the stock-selection strategy is represented by standard Brownian motion dB_t^i in normal times and by the random variable $\tilde{\omega}^i$ during crashes; the crash risk of the long-term trade is represented by a Poisson process dN_t with an intensity δ_N^i.

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

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

where r is the riskless rate, α^i is the manager's alpha, σ is the volatility of the security selection strategy in the absence of crashes, and y^i_{t+} is an i.i.d. standard normal drawn when a crash dN_t hits. The skilled manager generates excess returns on her selection portfolio, while the unskilled manager tries to select securities but has no skill, $\alpha^g > \alpha^b = 0$.

The price of the long-term trade evolves as,

$$dA_t = (r + \lambda)\, dt + \xi dN_t, \tag{2}$$

In normal times, this opportunity returns $\lambda > 0$ in excess of the riskless rate; when it crashes, it returns $\xi < 0$. We assume that both managers and investors know ξ. 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. When $E^g[dA] < r$ the optimal strategy resembles an event-driven style because the maximum expected return portfolio pays the carry cost $-\lambda dt$, but earns a large profit $(-\xi)$ when the event dN_t realizes. When $E^g[dA] > r$ the optimal portfolio resembles a volatility-insurance selling strategy which earns the carry λdt but loses when the event happens. We assume the unskilled manager underestimates the size or frequency of crashes and therefore perceives incorrectly that the volatility-insurance selling strategy is optimal, that is, $E^b[dA] \geq \lambda + r$. We assume that the idiosyncratic shocks in the model are manager-specific.

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

$$\begin{aligned} dR^i_{t+} &= dS^i_t + x^i_t(dA_t - rdt) \\ &= \left(r + \alpha^i + x^i_t\lambda\right) dt + \sigma dB^i_t + \left(x^i_t\xi + y^i_{t+}\right) dN_t, \end{aligned} \tag{3}$$

where we use dR^i_{t+} notation to describe the return process including the potential crash and dR^i_t to denote the return process before the potential crash; that is, if there is no crash at date t, $dR^i_{t+} = dR^i_t$. We assume that funds pay dividends to keep their sizes constant at one—we describe these dividend policies below—and so R^i_{t+} is the cumulative dollar gross return up to date t.

The manager invests $x^i_t \in [-1, 1]$ in the long-term trade, financing this position by buying or selling the riskless asset. 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] We assume that the unskilled manager behaves non-strategically and has always a long position in the long-term trade. The difference in the managers' positions in the long-term trade is therefore always weakly negative, $x^g - x^b \in [-2, 0]$.

Equation (3) is a stochastic integro-differential equation—these are distinct from stochastic differential equations in that they feature non-local movements. The presence of non-local movements compromises analytical tractability but it is essential for studying how endogenous liquidation risk alters the way managers value and invest in strategies with a long-term payoff profiles. These are strategies that likely pay poorly in the short-term but that have high expected returns.

2.1.1 Discussion

The assumption that the long-term trade is exposed to crash risk is a modeling device for capturing a trade-off between short- and long-term returns. When the impending crash is small, the manager faces no trade-off: the trade that maximizes expected returns is the same that maximizes short-term returns, that is, returns conditional on there not being a crash. However, when the impending crash is large or crashes frequent, the skilled manager must short the long-term trade and pay λ to profit from the eventual crash. In this case the manager has to decide between maximizing expected returns and

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

signaling skill by enhancing short-term returns. If the manager maximizes expected returns, investors may interpret her low short-term returns as evidence of lack of skill and redeem their investments.

The long-term trade 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' 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 Learning

Investors have a prior belief P_0 that the manager is skilled. Investors update their beliefs P_t as Bayesians by observing returns. They also form beliefs about skilled managers' strategies. Investors do not know the manager type, but they understand that the manager's optimal choice x_t depends on it.

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 long-term trade, they will not update their beliefs when there is a crash. If, however, they believe that the managers have different exposures, crashes are informative.

We use $E^I[z_t]$ to denote date t conditional expectations from the investors' perspective, that is,

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

$E^I[z_t] = E[z_t|\mathcal{F}_t^I]$ where is $\mathcal{F}_t^I$ is the filtration generated by the cumulative fund return R_t. We denote expectations conditional on a crash realization by $E_+^I[z_t] = E[z_t|\mathcal{F}_t^I, dN_t = 1]$. We omit the t subscripts from the expectations for notational convenience.

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 on convenience. We refer to the probability that the manager is skilled as the manager's *reputation*. We also define the difference and average operators for any manager type-specific random variable z^i,

$$\mathcal{D}(z) = z^g - z^b, \tag{4}$$

$$\mathcal{A}(z) = \frac{z^g + z^b}{2}. \tag{5}$$

Proposition 1 in the appendix shows that the log-likelihood ratio of investors' beliefs about skill evolves as

$$dp_{t+} = \frac{E^I[\mathcal{D}(dR_t)]}{\sigma^2}\left(dR_t^i - E^I[\mathcal{A}(dR_t)]\right) + \frac{E_+^I[\mathcal{D}(dR_{t+})]}{\omega^2}\left(dR_{t+}^i - E_+^I[\mathcal{A}(dR_{t+})]\right) dN_t. \tag{6}$$

If investors expect skilled managers to display better short-term performance, they infer from positive return surprises ($dR_t^i - E^I[\mathcal{A}(dR_t)] > 0$) that the manager is more likely skilled. Investors evaluate performance relative to their expectation about the average manager's performance, $-E_t^I[\mathcal{A}(dR_t)]$. The sensitivity of a manager's reputation to performance surprises depends on the informativeness of short-term performance; this informativeness is given by the difference between the skilled and unskilled managers' expected short-term performance scaled by return volatility, $\frac{\mathcal{D}(\alpha)+E^I[\mathcal{D}(x_t)]\lambda}{\sigma}$. Analogously, the informativeness of crash performance is given by the difference in the expected crash-event performance normalized by the crash-event idiosyncratic risk volatility, $\frac{E_+^I[\mathcal{D}(dR_{t+})]}{\omega}$. Manager reputation is more sensitive to short-term performance when investors expect skilled and unskilled managers to hold similar portfolios; it is more sensitive to crash-event performance when the expected portfolio differences are large.

A key aspect of the learning problem is that investors do not know managers' positions in the long-term trade. Investors therefore have to form beliefs about managers' choices to be able to interpret the performance signals; but this inference problem then in turn affects managers' portfolio choices. We highlight this mechanism by characterizing how the manager's portfolio choice affects the distribution of reputation. We can express the reputation dynamics in terms of the manager's portfolio choice and investors' beliefs,

$$var(dp_{t+})/dt = \left(\frac{E^I[\mathcal{D}(\alpha + x_t\lambda)]}{\sigma}\right)^2 + \delta \left(\frac{E^I[\mathcal{D}(x_t\xi)]}{\omega}\right)^2, \tag{7}$$

$$E[dp_{t+}]/dt = \frac{E^I[\mathcal{D}(\alpha + x_t\lambda)]}{\sigma^2}\left(\alpha^i + x_t^i\lambda - E^I[\mathcal{A}(\alpha + x_t\lambda)]\right) + \delta \frac{E^I[\mathcal{D}(x_t\xi)]}{\omega^2}\left(x_t^i - E^I[\mathcal{A}(x_t)]\right)\xi. \tag{8}$$

Investors' beliefs about the difference in portfolios, $E^I[\mathcal{D}(x_t)]$, have the opposite effects on the volatility of reputation in normal times and during crashes. As the expected difference becomes more negative, the volatility of normal-times reputation increases and that during crashes increases. The total volatility of reputation is U-shaped in the investors' perceived difference in portfolios. When the perceived difference is small, the low normal-times volatility dominates; but when the perceived difference is large, the crash-event volatility dominates. The expected growth rate of reputation similarly depends on the manager's actual portfolio choice and investors' beliefs about this choice. As the perceived difference in the portfolios grows wider, normal-times growth decreases and the crash-event growth increases. An increase in the actual portfolio choice, x^i, increases expected reputation growth in normal times and reduces it during a crash event. The dependency between overall reputation growth and the manager choice in turn strongly depends on investors' beliefs,

$$\frac{\partial E[dp_{t+}]/dt}{\partial x^i} = \frac{\mathcal{D}(\alpha)}{\sigma^2}\lambda + E^I[\mathcal{D}(x_t)]\left(\delta\left(\frac{\xi}{\omega}\right)^2 + \left(\frac{\lambda}{\sigma}\right)^2\right). \tag{9}$$

When investors expect no portfolio differences $E^I[x^g] = E^I[x^b] = 1$, reputation growth is strictly increasing in x^i and $x^i = 1$ maximizes reputation growth. If investors instead expect large differences, $E^I[\mathcal{D}(x_t)] = -2$, then reputation growth is strictly decreasing in x^i and $x^i = -1$ maximizes reputation growth. This result illustrates the strong strategic complementaries between investors' beliefs and manager behavior. Equation (9) implies that there is a threshold $\hat{x} = 1 - \frac{\mathcal{D}(\alpha)\lambda}{\delta\frac{\xi^2}{\omega^2}\sigma^2+\lambda^2}$ such that if $E^I[x^g] < \hat{x}$, then $x^i = -1$ maximizes reputation growth; if instead $E^I[x^g] > \hat{x}$, it is $x^i = 1$ that maximizes reputation growth. This threshold is increasing in the selection skill difference $\mathcal{D}(\alpha)$ between good and bad managers; it is decreasing in the ratio of systematic to idiosyncratic volatility during crashes, $|\xi|/\omega$.

2.3 The management contract

The contract between the manager and the investors specifies a management fee f and a performance fee m for returns in excess of the riskless rate. We initially assume that investors withdraw their funds and the manager is liquidated when the manager's reputation falls to a threshold $\underline{p}$. We later show that investors optimally follow such a liquidation policy when their holdings in the fund are liquid; at that time we also introduce and examine the consequences of lockup provisions.

The fund pays dividends to investors so that it always has one unit of capital under management.

These dividends therefore equal

$$dD^i_{t+} = rdt + (1-m)(dR^i_{t+} - (r+f)dt). \tag{10}$$

2.4 The manager problem

Both skilled and unskilled managers discount cash flows at the riskless rate r. By managing a fund, a manager earns the management fee and performance fee specified in the contract up to the moment she is liquidated. The manager's valuation therefore depends on both the contract terms and her reputation, which influences the probability of liquidation. The manager's reputation and compensation both depend on portfolio choices, and the manager chooses an investment policy x_t to maximize lifetime value:

$$G^i(p_t) \quad = \quad \max_{x^i_t} \mathrm{E}^i \left[\int_0^{\tau(\underline{p})} e^{-rt} \left[m\left(dR_{t+} - (f+r)dt\right) + f dt \right] \right], \tag{11}$$

which is the present value of the management and performance fees the manager earns up to the uncertain time $\tau(\underline{p})$ at which the fund is liquidated.

The value process $G(p_t)$ is a martingale for $t < \tau(\underline{p})$ that satisfies the following Hamilton-Jacobi-Bellman (HJB) equation:

$$\begin{aligned} rG^i(p_t) \quad &= \quad \max_{x^i_t} \Bigg\{ f + m(\alpha^i + x^i_t(\lambda + \delta\xi) - f) + G^i_p(p_t) \frac{E^I[\mathcal{D}(\alpha + x_t\lambda)]}{\sigma^2} \left(\alpha^i + x^i_t\lambda - E^I[\mathcal{A}(\alpha + x_t\lambda)]\right) \\ &\quad + \frac{1}{2} G^i_{pp}(p_t) \left(\frac{E^I[\mathcal{D}(\alpha + x_t\lambda)]}{\sigma} \right)^2 - \delta \mathrm{E}^i_t \left[G^i(p_{t+}) - G^i(p_t) \right] \Bigg\}, \end{aligned} \tag{12}$$

$$p_{t+} \quad = \quad p_t + \frac{E^I[\mathcal{D}(x_t\xi)]}{\omega^2} \left(x^i_t\xi - E^I[\mathcal{A}(x_t\xi)]\right) + \frac{E^I[\mathcal{D}(x_t\xi)]}{\omega} y^i_{t+}. \tag{13}$$

We normalize the manager's outside option to be equal to zero, $G^i(\underline{p}) = 0$; this is the lifetime utility that the manager earns after fund liquidation. Equation (12) provides intuition for the forces that shape the manager's decisions. On the left we have the time discounting of the value function. The first two terms on the right represent the manager's expected compensation: management fee and the expected performance fee during both normal times and crashes. The third term is the valuation effect of the expected change in the manager's reputation; reputation is a martingale only from the investors' vantage point. The second to the last term is a Jensen's inequality term. The concavity or convexity of the value function determines whether the manager favors or shuns reputation volatility. The last term represents the valuation effect of reputation changes due to a crash event.

Propositions 2 and 3 show that the manager's valuation, with some restrictions on the parameters, is increasing and concave in reputation. Intuitively, the manager's compensation depends on her reputation only through its impact on the odds of fund liquidation. Therefore, when the manager's reputation is high, further increases in reputation have little value as they do not materially alter the odds of fund liquidation. However, when the manager's reputation is low and close to the liquidation threshold, an increase in reputation will substantially reduce the probability of liquidation. We show next that the concavity of the manager's value function implies that the (skilled) manager's value increases more in reputation during normal times than during a crash. This result lies at the core of this paper's mechanism: a manager who faces a meaningful risk of being liquidated displays a short-term bias in investment decisions to build reputation.

We can examine the determinants of the skilled manager's portfolio choice by focusing on the manager's first-order condition with respect to her position in the long-term trade, x_t,

$$\underbrace{m(\lambda+\delta\xi)}_{\substack{\text{compensation}\\\text{incentives}}}+\overbrace{\underbrace{G_p^i(p_t)\frac{E^I[\mathcal{D}(\alpha+x_t\lambda)]}{\sigma^2}\lambda}_{\text{short-term}}+\underbrace{\delta\frac{E^I[\mathcal{D}(x_t)]\xi^2}{\omega^2}\mathrm{E}^i\Big[G_p^i(p_{t+})\Big]}_{\text{long-term}}}^{\text{reputation concerns}}. \tag{14}$$

If the manager's optimal position lies in the interior, $-1 < x_t < 1$, this first-order condition equals zero. Equation (14) 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, which we discuss below, the manager's portfolio will be at one of the extremes, at either $x_t^g = -1$ or 1.

Performance fee m 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 (14) is negative when the long-term strategy is profitable.

The second term in equation (14) represents short-term reputational concerns. This term is positive when investors view high short-term returns as positive signals about managerial skill, that is, $E^I[\mathcal{D}(\alpha+x_t\lambda)] > 0$. This short-term reputational effect gives the manager an incentive to take a long position in the long-term trade regardless of the size of the impending crash. The more informative short-term returns are about skill, the higher the manager's incentive to distort her portfolio.

The last term in equation (14) represents long-term reputational concerns. A manager's portfolio choice, through its effect on fund returns, 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 long-term trade even if it is not profitable to do so, she is concerned about what doing so would do to her reputation if there is a crash. These long-term concerns therefore incentivize the manager to maximize expected returns if investors expect a skilled manager outperform an unskilled manager during a crash, that is, $E^I[\mathcal{D}(x_t)\xi] > 0$.

The relative importances of the short- and long-term reputational concerns depend on the manager's horizon. This horizon appears in equation (14) as the difference in the slope of the value function with respect to reputation (G_p) before and after the crash. These slopes are close to each other when a manager has a long horizon; for a manager close to liquidation, however, the before-crash value function is much steeper as small increases in reputation greatly reduce in the probability of liquidation.

2.5 Equilibrium and feedback loops

The manager's portfolio choice in Equation (14) depends not only on expected returns on the long term opportunity $\lambda - \delta\xi$ and endogenous reputation concerns G_p, but also on investors' beliefs about the manager's portfolio choice through the terms $E^I[\mathcal{D}(x_t)]$ and $E^I[\mathcal{A}(x_t)]$. This feature is typical of learning environments where the content of the signal depends on the expected action and the action taken, in turn, depends on the signal it conveys. This feedback loop generates an amplification mechanisms and creates equilibrium multiplicity.

A manager faces a greater threat of liquidation when her reputation falls. This decrease gives the manager an incentive increase in x_t^g towards one to perform well in the short-term to gain reputation. Investors realize this shift in incentives and adjust upwards their beliefs about the gap between the short-term performance of skilled and unskilled managers and adjust downwards their beliefs about the gap between their long-term performance. This inference further increases the manager's incentive to perform well in the short term through the second term in Equation (14), increasing x_t^g even further.

Equilibrium multiplicity arises when reputational concerns G_p are large relative to compensation incentives. This multiplicity emerges through investors' self-fulfilling beliefs. If investors believe that the manager will not invest in the long-term trade, then the manager's incentive to invest in the long-term trade decreases because of the way investors will learn. The same mechanism applies to both managers' short- and long-term reputational concerns—that is, how investors learn from normal-times and crash-event performance. In both cases the variation in investors' beliefs change managers' incentives in the direction of investors' beliefs. This generates strategic complementarity between beliefs and choice, resulting in equilibrium multiplicity.

Before turning to equilibrium portfolio, we first state the formal definition of equilibrium in this economy. The equilibrium imposes consistency between the investors' beliefs and the actions of skilled managers. We solve a Markovian equilibrium where the state variable is the manager's reputation (P_t). Investors' learning can only depend on the manager's reputation and their beliefs about skilled managers' actions. Skilled managers' decisions depend on their reputation. We assume that investors and managers play the equilibrium that maximizes the fund expected returns. Because investors and managers can communicate, we focus on the equilibrium with the highest expected returns; this equilibrium features higher compensation for the manager and higher net distributions to investors.

Definition 1. ***Markovian equilibrium.*** *A Markovian equilibrium is given by skilled managers' portfolio* x_t^g*, investors' beliefs about this portfolio* $x_t^{g,I} = E^I[x_t^g]$*, and a law of motion for investors' beliefs such that:*

(1) Skilled managers' portfolio choices are optimal given investors' beliefs;

(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; and

(4) skilled managers' portfolio choices maximize expected returns subject to conditions (1) through (3).

We have thus far specified only that investors are Bayesians and that they liquidate a fund when the manager's reputation falls to some threshold. These two assumptions suffice for describing the key aspects of equilibrium in our economic environment. We later endogenize the investor liquidation decision to build on this equilibrium definition.

2.5.1 Equilibrium Portfolios

Proposition 4 describes the skilled manager's optimal portfolio choice. If the long-term trade is not profitable, the manager's optimal strategy is simply $x_t^g = 1$; in this case, the strategy that maximizes expected returns is the same that maximizes short-term returns and therefore best signals skill. When the long-term strategy is optimal, the optimal choice ranges from $x_t^g = -1$ to $x_t^g = 1$ depending on the strengths of the short- and long-term reputational effects. If the short-term reputational concerns are small enough, the manager maximizes long-term expected returns by choosing $x_t^g = -1$. 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 $x_t^g = 1$. There is also an intermediate region with $-1 < x_t^g < 1$; here, the marginal long-term reputational concerns exactly balance out the short-term concerns and payoff incentives.

Panel A of Figure 1 illustrates how the equilibrium portfolio choice depends on reputation when the long-term trade is profitable, that is, when the impending crash ξ is large or the crashes frequent. The "liquidation region" in the graph is the region where the manager's reputation is below the liquidation threshold $\underline{p}$.

If the manager's reputation is low but above the liquidation threshold, the manager has a strong incentive to maximize short-term performance. She thus invests $x_t^g = 1$ in the long-term trade despite

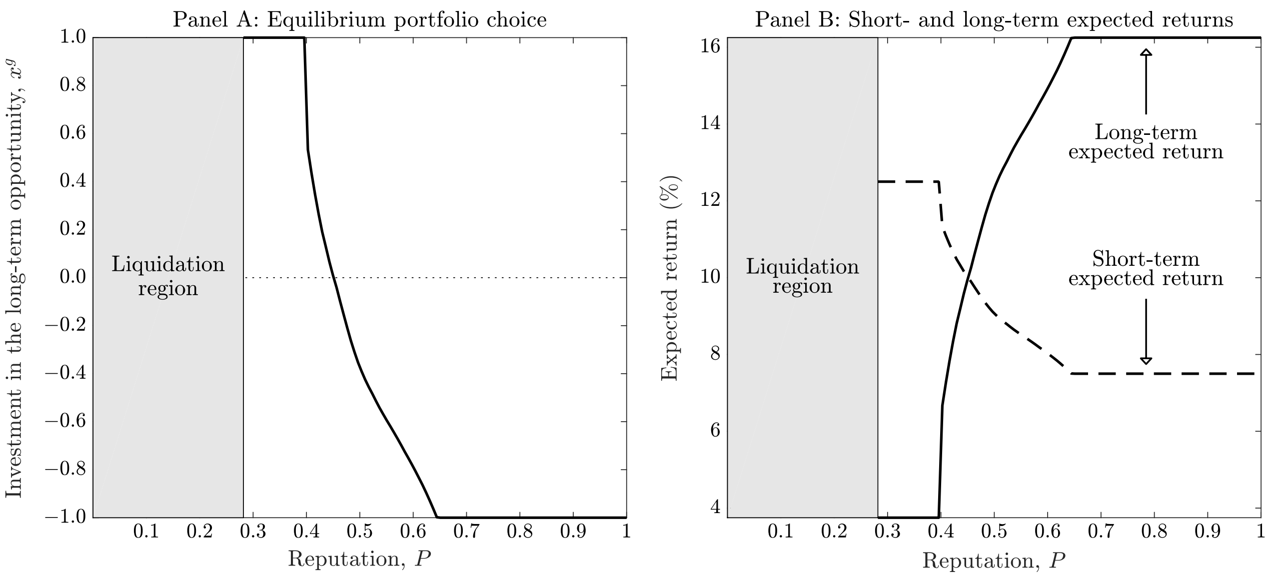

Figure 1: **Skilled manager's equilibrium portfolio choice and signaling.** This figure shows how the skilled manager's equilibrium portfolio depends on her reputation when the impending crash is large or the crashes frequent, $E^g[dA] < r$. The expected return-maximizing strategy pays the carry cost $-\lambda dt$, but earns a large profit $(-\xi)$ when there is a crash. Panel A shows the optimal portfolio x^g and Panel B plots the expected short- and long-term returns. The expected short-term return is the expected return conditional on there not being a crash. The expected long-term return accounts for the fact that crashes occur with intensity δ_N^g. The shaded area on the left denotes a manager whose reputation is too low to raise capital. As the manager's reputation increases, the manager does not need to improve her short-term returns to signal skill, and therefore invests in the strategy that maximizes expected (long-term) returns, $x^g = -1$.

knowing that she will suffer significant losses if there is a crash. This is the optimal choice because the manager is so close to the liquidation threshold; she cannot afford to maximize long-term expected returns because, in expectation, it will take too long for the long-term trade to pay off.

As the manager's reputation improves, her reputational concerns lessen, and she shifts towards the strategy that maximizes long-term expected returns. However, only when the manager's reputation is very high—here, if the probability that the manager is skilled is above 0.65—he follows the strategy that maximizes long-term expected returns, entirely ignoring short-term signaling issues.

Panel B of Figure 1 shows that reputational concerns limit arbitrage. In this figure, the expected short-term return is the expected return conditional there not being a crash; the expected long-term return is the (true) expected return that accounts for the fact that crashes occur with intensity δ_N^g. A skilled manager with a low reputation willingly forgoes expected returns to signal her ability to ensure survival. The gap between the maximum attainable expected return and the expected return that the manager delivers in the equilibrium measures limits to arbitrage. In Figure 1 the manager would obtain the maximum expected return of 16.25% by shorting the long-term trade. A manager who is just above the liquidation threshold, however, earns an expected return of just 3.75%. Such a manager trades against the long-term trade, making the gamble that there will not be a crash until she has built enough reputation to pursue the expected return-maximizing strategy.

This limits-to-arbitrage mechanism can be so severe that a fund completely unravels—that is, investors can be unwilling to invest in the fund no matter the manager's reputation. Proposition 5 shows that if the long-term trade is sufficiently profitable and crash-event returns are either very volatile or not volatile at all, that is, if $\omega \to 0$ or ∞, then the skilled manager's expected return is always negative when the manager is at the liquidation threshold. This result holds for any liquidation threshold $\underline{p}$, and so it leads to an unraveling that does not depend on the specifics of investors' preferences. If the skilled manager's expected return is negative at the liquidation threshold, then investors will be unwilling to invest at that reputation level. As a consequence, the liquidation threshold has to be somewhat higher. But because the result in Proposition 5 holds for any threshold level $\underline{p}$, there will be no reputation at which investors break even.

2.6 Lockup provisions

Investors and the fund manager can specify a lockup provision. The lockup expiration is a Poisson event with an intensity $\frac{1}{T}$. Investors who invest in a fund with a T-year lockup therefore expect the lockup to expire in T years. When the lockup expires, all investments become liquid and can be redeemed freely; before this date, all capital is locked up and cannot be redeemed. 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 trade is also a Poisson process—that is, its expected payoff also does not vary with the passage of time.

In Appendix A.2 we describe the full dynamic problem that the manager faces. Here we only highlight the differences relative to Equation (12), which gives the HJB equation in the absence of lockup provisions. The HJB equation is different depending on whether the fund is open (denoted by $l_t = 0$) or under an active lockup (denoted by $l_t = 1$),

$$rG^i(p_t, l_t = 0) = \ldots + \phi\Big[G(p_t, l_t = 1) - G(p_t, l_t = 0)\Big], \tag{15}$$

$$rG^i(p_t, l_t = 1) = \ldots + \frac{1}{T}\Big[G(p_t, l_t = 0) - G(p_t, l_t = 1)\Big], \tag{16}$$

where "…" represents the entire right hand side of Equation (12). When the lockup has expired, the valuation respects a boundary condition imposed by investors' liquidation policy so that $G(p_t, l_t = 0) = 0$ for all $p_t \leq \underline{p}_0$; here, we use a subscript 0 to denote the liquidation threshold in a fund with an inactive lockup. Equations (15) and (16) constitute a system of two coupled integro-differential equations. There are two equations because the system jumps between funds with active and expired lockups ($l_t = \{0, 1\}$).

The difference for a fund without an active lockup (see Equation (15)) is that it can transition into the active lockup state with intensity ϕ. We later show that these transition dynamics emerge endogenously from investors competing for managerial talent. The HJB equation changes in two ways for a fund with an active lockup (see Equation (16)). First, and more importantly, there is no liquidation boundary: the fund remains active for any reputation p_t. Second, the lockup expires with intensity $\frac{1}{T}$.

Skilled managers benefit from the lockups because they allow managers to take the long-term view.

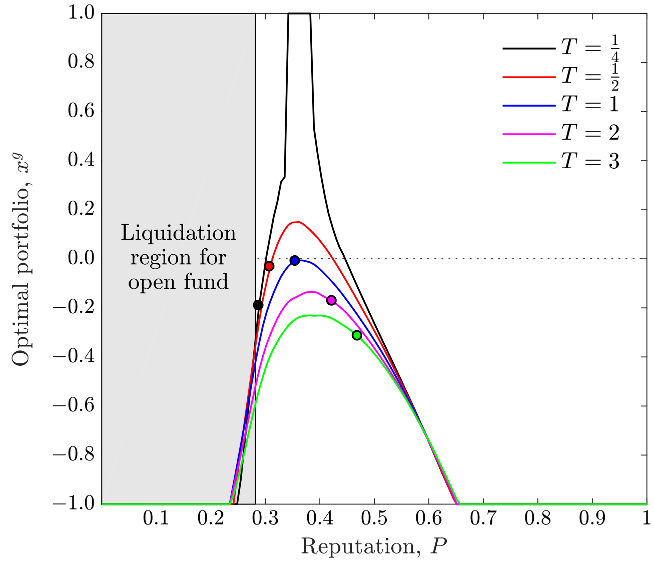

Figure 2: **Optimal portfolio choices with lockup provisions.** This figure shows how a skilled manager's optimal portfolio choice, x^g, depends on her reputation when the impending crash is so large that it is profitable to be short the long-term trade, $E^g[dA] < r$. The manager runs a fund with a lockup that ranges from three months (black line) to three years (green line). The shaded area signifies the liquidation threshold of an open-ended fund; the circles denote investment thresholds of the funds with lockup provisions; a manager is unable to raise capital if her reputation lies below these thresholds. Because a manager's reputation can fall below the investment threshold while the lockup is in effect, the lines denoting optimal portfolio choices extend to the left of the investment thresholds.

Managers do not have to worry as much 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 trade pays off while the lockup is active, the 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 2 illustrates the benefit of lockups. In this figure, skilled managers of different reputations manage funds with a lockup ranging from three months to three years. The gray area denotes the liquidation threshold of an open-ended fund. The circles denote *investment* thresholds; a manager is unable to raise capital if her reputation lies below these thresholds. After opening a fund, a manager's reputation can fall below the investment threshold, and so the lines denoting optimal portfolio choices extend to the left of the investment thresholds. 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 trade. 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 trade without worrying about the reputational cost of low short-term returns. As reputations fall, however, managers reduce their positions in the long-term trade. 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, she invests $x_t^g = 1$ to maximize her short-term returns. That is, she sacrifices expected returns to build reputation quickly. The portfolio choices of managers who run funds with lockup provisions begin to fall back down; that is, they again begin to trade the long-term trade 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 long-term trade and hoping that it pays off in time.

The change in manager willingness to bet in the long-term opportunity can be trace out to the shape of the manager value function G. While in an open-ended fund G is concave and G_p is maximum at the liquidation, a manager with a lockup provision has a value function that is S-shaped. This S-shape induces state-dependent variation in the manager's aversion to reputation shocks. Managers close to the liquidation threshold value small local increases in reputation the most because when slight above or slight below liquidation, small increases in reputation lead to the largest reduction in the probability of liquidation.

2.7 Investors and the market for skill

Our analysis above assumes that the manager faces an exogenous liquidation threshold $\underline{p}$ and that a fund with an expired lockup reinstates the lockup with probability ϕ. We close the model by introducing the investor's decision to invest in and to redeem from a fund. The objective is twofold. First, by modeling the redemption decision we can study portfolio distortions. We can measure the intensive margin of limits to arbitrage and quantify the feedback from the changes in the optimal liquidation threshold into the probability of liquidation. Second, by modeling the investment decision, we can examine how competition for managerial skill erodes the rents of investing with a manager who turns out to be better than expected, but locks investors in with bad managers. We study how this asymmetry impacts the thresholds at which investors are willing to invest and how much they are willing to pay in fees, thereby giving us a complete theory of lockup maturity choice.

Investors value a fund at either $V(p_t, l_t = 1)$ or $V(p_t, l_t = 0)$ depending on its lockup status. We begin by taking these valuations as given and model the market for skill. The market for skill functions as follows. A unit mass of potential investors gets matched with a fund with intensity ϕ. If they value a fund with an active lockup at a premium relative to the fund's net asset value, $V(p_t, l_t = 1) \geq 1$, they bid $F(p_t, l_t = 1) = V(p_t, l_t = 1) - 1$. Because investors compete against each other to invest, they transfer the entire surplus at the time of investment to the manager. The manager accepts this bid, returns capital to the old investors, accepts capital from the new investors, and changes the new investors a one-time load fee of $F(p_t, l_t = 1)$. If investors value a fund with an active lockup at a discount, $V(p_t, l_t = 1) < 1$, the outside investors walk away and nothing changes for the manager or the existing investors. If the outside investors encounter a fund without an active lockup and they value the fund at a premium only if they will not face a lockup, $V(p_t, l_t = 0) > 1 > V(p_t, l_t = 1)$, the outside investors bid $F(p_t,0) = V(p_t,0) - 1$ for an investment in the fund with an inactive lockup contract. The manager accepts the bid, continues to manage a fund without an active lockup, and earns a one-time fee of $F(p_t,0)$.

This model captures the fact that a manager who has a fund with an active lockup and who performs well can attract new investors, thereby diluting the value of the old investors' investments. We model this dilution as being complete by having managers return old investors their entire capital. In this setting managers with an inactive lockup can attract new investors who are willing to be locked up; this captures a key aspect of how lockups in practice—managers get to renew their lockups as long they are performing well enough.

The introduction of the market for skill again changes the manager HJB:

$$rG^i(p_t,0) = \ldots + \phi I_{(p_t \geq \underline{p}_1)} \Big[G(p_t,1) - G(p_t,0) + F(p_t,1) \Big] + \phi I_{(\underline{p}_1 > p_t > \underline{p}_0)} F(p_t,0), \quad (17)$$

$$rG^i(p_t,1) = \ldots + \frac{1}{T} \Big[G(p_t,0) - G(p_t,1) \Big] + \phi I_{(p_t > \underline{p}_1)} F(p_t,1), \quad (18)$$

where "..." represents the entire right hand side of Equation (12) and $I_{(z)}$ is an indicator function that returns 1 if z is true and 0 otherwise. The probability that a fund transitions from having an inactive lockup to having an active lockup depends on both the intensity ϕ—which controls the competitiveness of the market for skill—and the manager's reputation. A manager's reputation has to exceed the investment thresholds $\underline{p}_1$ and $\underline{p}_0$ at which investors just break even for investments into funds with active and inactive lockups. The manager also earns the load fee F every time new investors come in. This fee, in turn, depends on whether the new investors invest in the fund with an active or inactive lockup contract, and this decision is a function of the manager's reputation and the fund's current lockup status.

We can now solve for investor valuations. Risk-neutral investors discount cash flows at the rate $\rho(T) = r + \nu T$, where ν represents an illiquidity premium; the discount rate increases in the length of time that the investor's capital is tied to an investment. All investors value managers the same way. Investors value a fund at either $V(p_t, L_t = 1)$ or $V(p_t, l_t = 0)$ depending on its lockup status,

$$
\begin{aligned}
V(p_t, 1) &= E^I\left[\int_0^{\tilde{\tau}_1} e^{-\rho(T)t} dD_t + e^{-\rho(T)\tilde{\tau}_1}\left(I_{(\tau_c=\tilde{\tau}_1)} + I_{(\tau_e=\tilde{\tau}_1)} V(p_{t+\tilde{\tau}_1}, 0)\right)\right], \qquad (19)\\
V(p_t, 0) &= E^I\left[\int_0^{\tilde{\tau}_0} e^{-\rho(T)t} dD_t + e^{-\rho(T)\tilde{\tau}_0}\right], \qquad (20)
\end{aligned}
$$

where $\tilde{\tau}_1$ is the time at which new investors outbid and replace old investors (τ_c) or the lockup expires (τ_e), whichever comes first. 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. The stopping time $\tilde{\tau}_0$ for an open fund therefore is the time when new investors outbid old investors (τ_c) or the manager's reputation falls below the liquidation threshold $\underline{p}_0$, whichever comes first.

We solve for the investor value function by solving jointly for the pair of value functions $V(p_t, 1)$ and $V(p_t, 0)$, and investment thresholds $\underline{p}_1$ and $\underline{p}_0$ that satisfy the pair of HJBs associated with Equations (19) and (20) (see Appendix A.2) with three boundary conditions: (1) $V(\underline{p}_0, 0) = 1$, (2) $V_p(\underline{p}_0, 0) = 0$, and (3) $V(\underline{p}_1, 1) = 1$. The first and third conditions are "value matching" conditions; the investor must be indifferent between one dollar inside or outside the fund at the threshold. The second condition is a "smooth-pasting" condition; when a fund is open ended, the investor has the option of delaying liquidation. This flexibility implies that the value function has a zero slope at the optimum threshold.[4] There is no smooth-pasting condition for the optimal decision to invest in the fund because investors cannot delay investment in the fund.

[4]See Dixit (1993) for a discussion of smooth pasting.

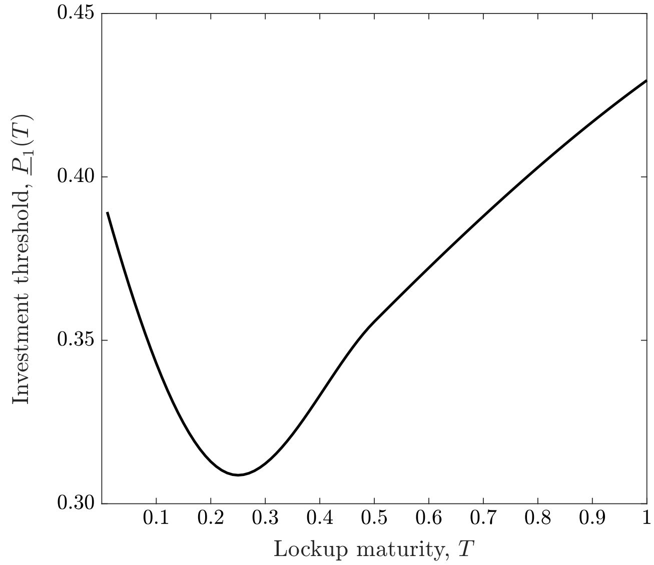

Figure 3: **Investment threshold as a function of lockup maturity.** A manager offers a fund to investors with a lockup maturity that ranges from one month to a year, and investors decide whether to provide capital. This figure plots the investment threshold as a function of lockup maturity. Investors are unwilling to give capital to a manager whose reputation lies below the investment threshold.

Proposition 6 shows that the investment threshold is increasing in the arrival rate of new investors. 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 low-reputation manager. 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. Figure 3 shows that this effect can be so strong that the investment threshold can decrease with lockup maturity; a manager with a reputation of, say, $P_t = 0.35$ would be unable to raise capital for an open-ended fund, but the same manager would be able to raise capital for a fund with a three-month lockup. This effect arises when investors expect managers with short-term contracts to distort their portfolios more and therefore have lower expected returns than managers with slightly longer-term contracts.

Lockups are costly to investors because they prevent investors from withdrawing capital in response to future bad news while, at the same time, investors cannot fully capture the rents of future good news. It is important to note, however, that it is exactly this restriction on investors' ability to liquidate after bad news that provides the manager the incentive to bet aggressively on the long-term trade.

2.7.1 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 perfectly 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 and bid away 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.

3 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	10%	Agarwal et al. (2009)
Crash volatility	ω	7.5%	
Long-term long-term trade			
Carry	λ	2.5%	Jurek and Stafford (2015)
Crash intensity	δ	0.5	
E[r] of the long-term trade	$-(\lambda+\delta\xi)$	5%	Estimated
Security selection abilities			
Skilled manager's alpha	α^g	10%	Match the average lockup
Unskilled manager's alpha	α^b	0%	premium of 4%
Outside offers			
Arrival rate	ϕ	1	Berk and Green (2004) and Aragon (2007)
Illiquidity cost	ν	1%	

3.1 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 set the crash volatility 7.5% which is the volatility experienced in a seven month period, $\omega = \sqrt{7/12} \times 10\% \approx 7.5\%$.

We choose the carry of the long-term trade (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% to this total. Based on these number we set the carry equal to $\lambda = 2.5\%$.

We assume a crash intensity of six months. Based on a calibration, which we discuss below, we set the expected return of the long-term trade to 5%. Together with λ and the crash intensity, the size of a crash is then $\xi = -15\%$. This is a modest crash in a risk-neutral world. Jurek and Stafford (2015), for example, estimate that the average hedge fund earned a return of -20% in 2008.

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^g because it is about the risk of being tied up with an unskilled manager who earns low returns. We set $\alpha^g = 10\%$ and $\alpha^b = 0\%$ which, in the model, generate a lockup premium of 3.8%.

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

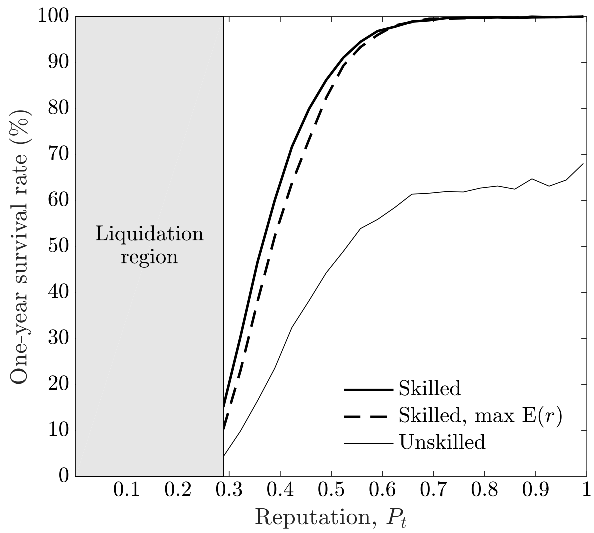

Figure 4: **Survival of skilled and unskilled managers.** This figure shows the one-year survival rates for skilled and unskilled managers of different initial reputations. The thick solid line represents skilled managers who make equilibrium portfolio choices. The thick dashed line represents skilled 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 $x_t = 1$.

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 $\phi = 1$. This choice implies that, on average, positive news are diluted in a year. We assume that investors have a preference for liquidity, and set the illiquidity premium at 1% per year.

3.2 Reputation, survival, and optimal portfolio choice

Managers who are concerned about their reputation trade the long-term trade cautiously; by distorting their portfolios, they can improve their chances of survival. Figure 4 shows how managers enhance their changes 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, the unlucky ones get liquidated. Because investors learn from returns, they sometimes mistake an unlucky skilled manager for an unskilled manager. This risk increases when we get closer to the liquidation threshold.

The thick solid line in Figure 4 shows the survival rates of managers who make optimal choices. This choice maximizes the manager's own valuation, which depends crucially on survival and 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 4 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 trade typically does not pay off in the short run, and so they instead 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, she does not need to distort her behavior to ensure survival.

Figure 5 shows how managers alter their behavior as they build reputation. In this analysis, we examine a manager with a reputation of 0.5. This manager distorts her portfolio moderately to signal skill. Instead of choosing $x_t = -1$ to maximize long-term expected returns when the impending crash is large, she chooses a portfolio of $x_t = -0.39$ (see Figure 1).

We now track this manager over time. Even when following the equilibrium strategy, 13 percent of these managers get liquidated over the first year. We also report the survival rates of unskilled managers.

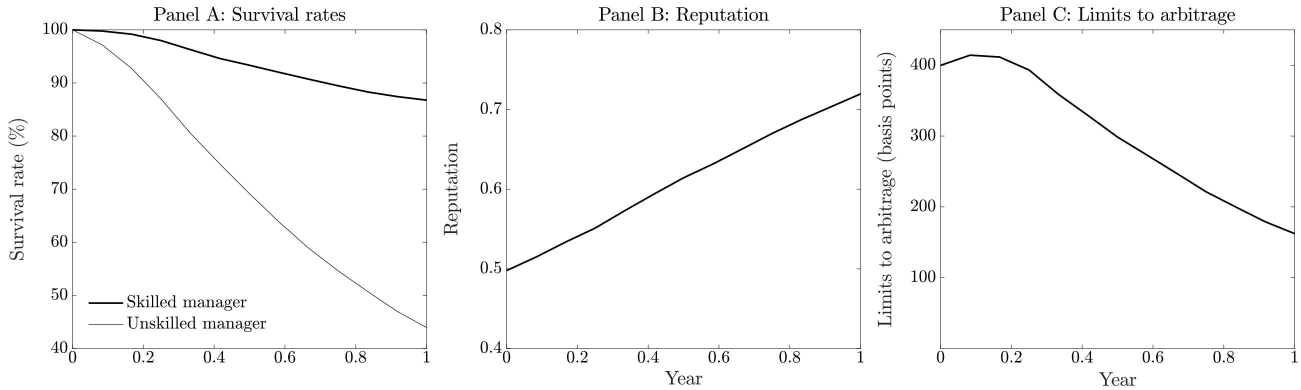

Figure 5: **Reputation and limits to arbitrage over time.** This figure simulates a year of data for skilled and unskilled managers. Each manager has an initial reputation of $P_t = 0.5$. Panel A plots the survival rates of these managers. The thick line represents skilled managers who make equilibrium portfolio choices. The thin solid line represents unskilled managers who always maximize short-term returns by choosing $x_t = 1$. 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 she obtains in equilibrium.

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 5 shows that, conditional on survival, skilled managers' reputations improve. If the manager survives for a year, her reputation has increased from 0.5 to 0.72. This increase reflects 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 5 shows the average limits to arbitrage, which is the gap between the maximum available expected return and the equilibrium expected return. 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 trade more aggressively as her reputation improves. This is the optimal choice; as the manager becomes less concerned about survival, the gain from higher returns offsets the reputational cost of lower short-term returns. In Panel C, the skilled manager initially forgoes 4% per year in returns to signal skill; this amount declines to 1.6% by the end of the first year.[5]

3.3 Signaling and commitment

In the model skilled managers sometimes deliberately destroy value to signal skill. Instead of earning an expected return of 16.25%, a skilled manager earns a return of just 3.75% in equilibrium; the threat of liquidation induces her to distort her 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 liquidation threshold is $\underline{P} = 0.28$, 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

[5]Limits to arbitrage initially increase in Panel C of Figure 5 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 long-term trade (see Figure 1), the average limits to arbitrage initially increases before beginning to decline.

to 0.19. That is, investors would now be willing to give capital to managers whose probability of being skilled is just one-fifth. 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 her gross return. The change in the liquidation threshold from 0.19 (with commitment) to 0.28 (without commitment) corresponds to investors only considering managers who can deliver 2.0% higher expected gross returns.

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

Skilled managers distort their portfolios because returns carry information that investors use to update their beliefs. Investors learn from both short- and long-term returns. The informativeness of these returns therefore determines the extent to which managers alter their behavior. In Figure 6, 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 normal-times and crash-event returns.

Panel A of Figure 6 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 or the crashes frequent so that the expected return-maximizing portfolio invests $x_t = -1$ in the long-term trade. 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 $x_t = 1$) is 0.35. When short-term returns become more informative, this threshold increases to 0.48. All managers below this level view the short-term

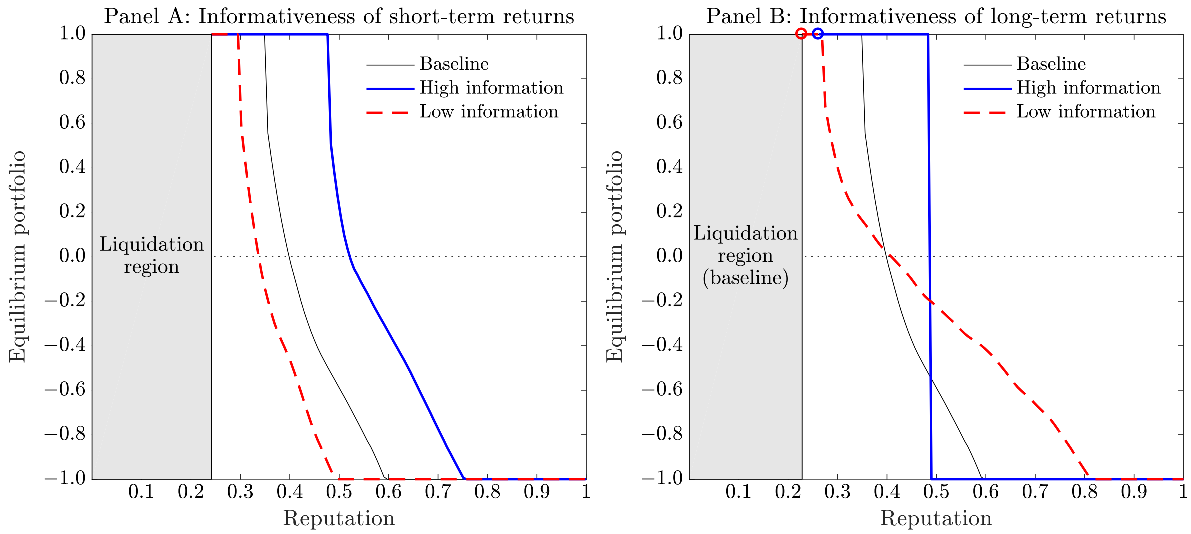

Figure 6: **Optimal portfolio choices when short- and long-term returns are more or less informative.** This figure shows how the skilled manager's equilibrium portfolio depends on her reputation when the impending crash is large or the crashes frequent, $E^g[dA] < r$. The skilled manager runs a fund 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 thresholds are (almost) the same in Panel A; Panel B plots the lowest threshold (the one in the low-information environment) and uses circles to indicate the liquidation thresholds corresponding to the baseline and high-information environments.

reputational cost of trading the long-term trade too high.

Panel B modifies investors' ability to learn from crashes. Long-term reputational concerns do not matter when when learning during crashes is either perfect of fully absent. At these extremes, manager faces almost no long-term reputational concerns. When learning is perfect, the manager knows that as long as her expected crash performance is infinitesimally better than that of the unskilled manager, investors will learn that she is skilled. At the other extreme, when crash returns carry little information, the manager's portfolio choice today does not affect her expected after-crash reputation. Long-term reputational concerns therefore matter only between these two extremes.

Panel B of Figure 6 illustrates the relations between reputation, the informativeness of crash-event returns, and the limits to arbitrage. When crashes carry more information about skill—this is the high-information environment in Panel B—all managers with reputations below 0.48 maximally distort their portfolios to earn high short-term returns to signal skill. In the low-information environment, only managers with reputations below 0.27 do so. That is, when crashes carry more information about skill, low-reputation managers are less eager to distort their portfolios. These managers are better off my betting that a crash occurs, thereby inducing a large positive shock to their reputations.

3.5 Maturity choice

We now examine the lockup maturity decision. Because investors and managers can bargain over the manager compensation, we can reduce this analysis to a measurement of the total value associated with a lockup contract. Although managers vary in their initial reputations p_0, every manager who is about to start a new fund believes that she is skilled. The manager first writes a contract that specifies the lockup maturity T, and offers this contract to investors. She chooses the lockup maturity to maximize her valuation,

$$\begin{aligned} T^*(p_0) &= \arg\max_T \left\{ F(p_0, l_0 = 1|T) + G^g(p_0, l_0 = 1|T) \right\}, \\ \text{subject to} \quad & p_0 \geq \underline{p}_1(T), \end{aligned} \tag{21}$$

where the constraint $p_0 \geq \underline{p}_1(T)$ verifies that the investment threshold associated with the proposed lockup maturity is low enough to attract capital. After starting a fund, a skilled manager makes her decisions according to equation (11) to maximize her value. An unskilled manager behaves non-strategically and earns the carry from the long-term trade—these managers can be viewed as maintaining their belief that they are skilled and consistently underestimating the size or frequency of the crashes.

The optimal maturity choice trades off the benefit of a reduction in the limits to arbitrage and the costs associated with entrenchment and illiquidity. The total value of the fund $F(\cdot) + G(\cdot)$ reflects the net benefits of an increase in lockup contract maturity; this is what the manager maximizes when choosing which maturity to offer to investors.

Figure 7 shows that the total fund value is strongly increasing in lockup maturity among low-reputation managers, but less so among high-reputation managers. Intuitively, low-reputation managers are close to the liquidation threshold and therefore gain the most from the extra time a longer lockup provides. At the same time, longer lockups are also the most costly to the investors of low-reputation managers because of the higher changes that it will entrench a bad manager. These considerations imply that it is the feasibility of a lockup that entirely drives the lockup decision among low-reputation managers; that is, the manager chooses the longest maturity that investors will accept. Among high-reputation managers, who are close to indifferent across different maturities, lockup maturity is an interior choice. These managers already have a long effective horizon; they would need to be exceptionally unlucky to be liquidated given their high reputation. We see this pattern in Figure 7 which shows the optimal lockup maturity (Panel A) and the valuation gain of the optimal contract relative to the shortest feasible contract (Panel B). These gains are large among low-reputation managers but fairly small for high-reputation managers.

As a manager's reputation improves, 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. Highly reputable managers would benefit little from the lockups. However, without additional restrictions, other things equal, they would still weakly prefer long-term contracts. Because high-reputation managers are almost indifferent between all contract maturities, even a small cost pushes them to favor short-term contracts.

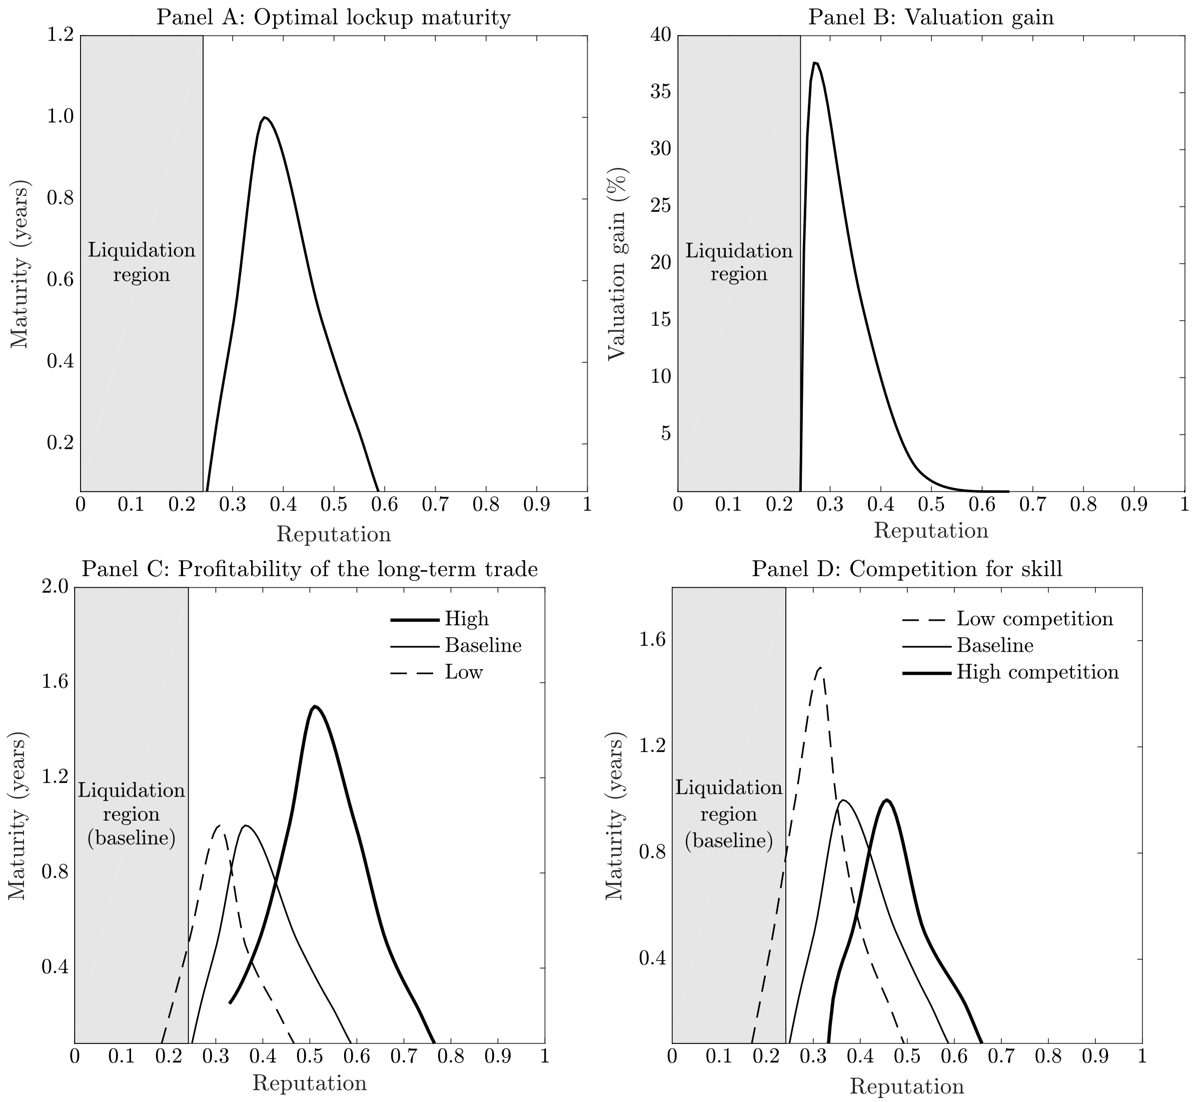

Figure 7: **Optimal lockup maturity: Valuation gains, profitability of the long-term trade, and competition for skill.** This figure shows the optimal lockup (Panels A, C, and D) and the valuation gain (Panel D) as a function of manager's reputation. A manager chooses the optimal lockup to maximize her valuation at the time she 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 profitability of the long-term trade; its expected return is 3% (low), 5% (baseline) and 10% (high). Panel D varies the competition for skill by changing the arrival rate of outside investors. The arrival rates are once in two years (low competition), once a year (baseline), and once in six months (high competition). In Panels C and D, the liquidation threshold depends on the parameters. The shaded area signifies the liquidation threshold in the baseline case.

Panels C and D show how the optimal lockup maturity responds to changes in the profitability of the long-term trade and the competitiveness of the market for skill. The expected return of the long-term trade is an important driver of the optimal maturity. In the baseline calibration, managers can earn an additional 5% in returns by trading the long-term trade. If this trade were to turn unprofitable, the signaling problem would disappear, and the lockup provision would no longer serve any purpose. Panel C shows that when the long-term is very profitable, the signaling problem worsens. Investors are willing to provide capital only to managers with reputations above 0.33; moreover, at this threshold, the optimal contract already specifies a lockup period of one quarter. This discontinuity is reflects the same mechanism as that highlighted in Figure 3; investors expect managers with short-term contracts to lower their returns so substantially by distorting their portfolios, that they find it optimal to start with a lockup at the threshold.

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 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 year to once in two years, a manager who was previously just able to raise capital with a reputation of $P_0 = 0.24$ can now get a contract with a nine-month lockup.

4 Implied distribution of skill

Optimal lockup maturities and survival rates depend on managers' reputations and the composition of manager skill. 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. Furthermore, maturity choice also reveals the profitability of the long-term trade; managers with more profitable long-term investment opportunities stand to gain more from negotiating longer contracts. In this section, we use these relations between reputations, the profitability of the long-term trade, optimal contracts, and survival rates to back out the implied distribution of reputation.

In Table 2 we illustrate the information contained in the distribution of lockup maturities. We estimate the distribution of skill as follows. We assume that the distribution of managers' (log-likelihoods of) reputations before they open a fund is Beta-distributed. We denote this distribution by $B(a, b)$. We then vary the profitability of the long-term trade by varying the crash size ξ.

For each a triplet $\{a, b, \xi\}$, we draw a large number of managers from the distribution $B(a, b)$, have those with sufficiently high reputations raise capital and start funds with optimal lockups, and simulate three years of data. We record the distribution of optimal lockup maturities and the funds' survival rates. We repeat these simulations to find a, b, and ξ by matching, between the simulations and the data, this distribution of lockups and the survival rates.

We report the results in Table 2. We report the estimates using the parameters listed in Table 1 ("baseline") as well as an alternative parametrization that lowers the illiquidity premium ν from 100 bps to 50 bps ("low illiquidity premium"). For both illiquidity premiums the model can perfectly match the first-year attrition ratio, which is around 18% in data (Brown, Goetzmann, and Park 2001), and the distribution of lockup maturities, which are 30%, 67%, and 3% for 0- to 3-month, 3- to 12-month, and longer-than-12-month maturities.[6] In the case of high (low) illiquidity premium, the estimated average reputation is 0.54 (0.49) and the profitability of the long-term trade is 5% (3%) per year.

With these estimates on hand, we can recover investment distortions from the distribution of lockup maturities. In the data lockup maturities are relatively short. When we assume that investors require only small compensation for illiquidity, our estimates imply that the investment distortions are small.

[6] We take the empirical distribution of lockup maturities from the Hedge Fund Research (HFR) database. Three 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.

Managers who open a new fund forgo on average 40 basis points per year to signal skill; this distortion. In the baseline specification, in which investors require a higher illiquidity premium, the model implies substantially higher limits to arbitrage. The average new fund manager foregoes 150 basis points per year in returns. These distortions dissipate over time as the reputations of the managers near the liquidation threshold either deteriorate (and the managers get liquidated) or improve. Among managers who survive a year, the average distortion falls from 150 to 80 basis points; in the low-liquidity premium estimation, the distortion falls from 40 to 20 basis points.

The last column of Table 2 illustrates model mechanics and the information content of lockup maturities by assuming that the distribution of lockup maturities in the data is substantially (and counterfactually) longer. Specifically, we assume that 23% of the funds have lockups with at least one-year maturities, and the share of funds with 3-to-12-month lockups is 47% instead of 67%. The estimate of the profitability of the long-term trade now increases to 6%—the longer lockup choices imply that managers must benefit greatly from longer lockup maturities, which the model interprets as a more profitable long-term trade.

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 the contract maturity choice 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

Table 2: Estimation

This table reports the moments we match between the model and the data, the estimated distribution of skill, and the implied distribution of reputation and the limits to arbitrage. We estimate the model using the parameters listed in Table 1 ("baseline") and using the same parameters except with the illiquidity premium lowered from 1% per year to 0.5% per year ("low illiquidity premium"). The last column assumes that the distribution of lockup maturities is substantially (and counterfactually) longer than what it is in the data. We estimate the first-year attrition rate and the distribution of lockup maturities from the Hedge Fund Research (HFR) database.

Moment or parameter	Data	Baseline	Low illiquidity premium	Counterfactually long lockups
Illiquidity premium		1%	0.5%	1%
			Moments	
First-year attrition rate	18%	18%	18%	17%
Distribution of lockups				
Less than 3 months	30%	30%	30%	30%
Between 3 and 12 months	67%	67%	67%	47%
At least 12 months	3%	3%	3%	23%
			Estimated distribution of skill, $B(a,b)$	
a		102.14	92.58	24.78
b		87.73	95.98	18.67
$E(r)$ of the long-term trade		5%	3%	6%
			Implied moments	
Average reputation ($P \in [0,1]$)		0.54	0.49	0.57
Reputation standard deviation		0.04	0.04	0.08
Limits to arbitrage				
New manager		1.46%	0.27%	2.10%
After one year		0.80%	0.31%	1.02%

and incubation biases, survivorship bias, liquidation bias, and self-reporting bias.[7] Using a structural

[7]See, 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 Ravenschaft (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.

model to estimate the distribution of skill is complementary to estimating alphas from these hedge fund databases.

5 Conclusions

We present a model in which managers improve their short-term returns at the expense of expected long-term returns to signal to investors that they are skilled. 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. This friction can generate large limits to arbitrage. Managers do not exploit the long-term investment opportunity to its full extent because they are concerned about investors interpreting the resultant low short-run returns as evidence of lack of skill. Limits to arbitrage is a robust feature of an environment in which investors are uncertain about managerial skill, and managers have to decide how aggressively to bet on the long-term trade.

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 get exposure to the long-term trade 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 investment? How informative is short-term performance about the managerial skill? How competitive is the market for skill? How costly is it to be stuck with an unskilled manager? 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, it applies to other contracts as well, such as those pertaining to CEO compensation. Whenever there is uncertainty about managerial skill and when profitable long-term opportunities may be costly in the short run, the insights of our model apply.

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.

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.

A Appendix

A.1 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*

$$dp_{t+} = \frac{E^I[\mathcal{D}(dR_t)]}{\sigma^2}\left(dR_t^i - E^I[\mathcal{A}(dR_t)]\right) + \frac{E_+^I[\mathcal{D}(dR_{t+})]}{\omega^2}\left(dR_{t+}^i - E_+^I[\mathcal{A}(dR_{t+})]\right)dN_t. \tag{22}$$

In equation (22)*,* dR_t *is the instantaneous non-crash return,* dR_{t+} *is the crash-event return,* $\frac{E^I[\mathcal{D}(dR_t)]}{\sigma}$ *and* $\frac{E_+^I[\mathcal{D}(dR_{t+})]}{\omega}$ *are the signal-to-noise ratios of short-term performance and crash performance, and* $E^I[\mathcal{A}(dR_t)]$ *and* $E_+^I[\mathcal{A}(dR_{t+})]$ *are investors' expectations about short-term and crash performance in log-likelihood space:*

$$\begin{aligned}
E^I[\mathcal{D}(dR_t)] &= (\alpha^g + E^I[x_t^g]\lambda) - (\alpha^b + \lambda), \\
E_+^I[\mathcal{D}(dR_{t+})] &= -\xi(E^I[x_t^g] - 1), \\
E^I[\mathcal{A}(dR_t)] &= \frac{(\alpha^g + E^I[x_t^g]\lambda) + (\alpha^b + \lambda)}{2}, \text{ and} \\
E_+^I[\mathcal{A}(dR_{t+})] &= -\xi\frac{E^I[x_t^g] + 1}{2}.
\end{aligned}$$

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 investors 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

$$P_{t+\Delta} = \frac{e^{-\frac{(R_{t,\Delta}-\mu_t^g\Delta)^2}{2\sigma^2\Delta}}P_t}{e^{-\frac{(R_{t,\Delta}-\mu_t^g\Delta)^2}{2\sigma^2\Delta}}P_t + e^{-\frac{(R_{t,\Delta}-\mu_t^b\Delta)^2}{2\sigma^2\Delta}}(1-P_t)}. \tag{23}$$

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

$$P_{t+\Delta} = \log\left(\frac{e^{-\frac{(R_{t,\Delta}-\mu_t^g\Delta)^2}{2\sigma^2\Delta}}P_t}{e^{-\frac{(R_{t,\Delta}-\mu_t^b\Delta)^2}{2\sigma^2\Delta}}(1-P_t)}\right) = P_t + \frac{\mu_t^g-\mu_t^b}{\sigma} \times \frac{1}{\sigma\Delta}\left(R_{t,\Delta} - \frac{\mu_t^g+\mu_t^b}{2}\Delta\right). \tag{24}$$

Investors' expectations about short-term performance depend on their beliefs about the skilled manager's portfolio choice, $E^I[x_t^g]$. The log-likelihood process of perceived skill, in absence of crashes, therefore evolves as given by the first term of Equation (22).

When there is a crash, the manager experiences a large return. Because the unskilled manager is always long the arbitrage opportunity, she loses $-\xi_t$ when the crash hits. The skilled manager's return depends on x_t^g. Investors again need to distinguish between two normal distributions that have different means but equal variances; the return volatility associated with the crash is ω. Applying Equation (24) to the crash return, we obtain the second term of Equation (22). □

Proposition 2. ***A manager's valuation increases in reputation.*** *Let $\phi = 0$, $\underline{G} = 0$, $f > 0$, and $x_t^g = E^I[x_t^g] = x_0$, and fix the liquidation threshold $\underline{p}$.The manager's valuation then increases in her reputation, $G_p > 0$.*

Proof. If no outside offers are made ($\phi = 0$) and the portfolio is fixed, 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, $\Pr(p_{t+\Delta} < \underline{p}|p_1) < \Pr(p_{t+\Delta} < \underline{p}|p_2)$.

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 $\underline{p}$, thus immediately follows that the probability of fund liquidation is increasing in probability of hitting $\underline{p}$. 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. 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 ($G_p > 0$).

□

Proposition 3. ***A manager's valuation is concave in reputation.*** *Let $\phi = 0$, $\underline{G} = 0$, $f > 0$, $x_t^g = x_0$, and $\omega \to \infty$, and fix the liquidation threshold $\underline{p}$ for the two lockup and crash states as well as the manager's portfolio. Then, A skilled manager who manages an open-ended fund is risk-averse with respect to reputation risk: $p_t \geq \underline{p} \Rightarrow G_{pp} < 0$.*

Proof. When $\omega \to \infty$ there is not learning during crashes. The results hold more generally, but the "no learning in crashes" case is convenient as the manager HJB becomes one ODE with constant coefficients, with boundary condition $G(p_t) = 0$ for $p_t \leq \underline{p}$. The constructive proof is almost identical in the opposite extreme, that is, if learning perfect so that a manager's reputation jumps to $p_t = \infty$ the first time a crash arrives. The ODE is given by

$$0 = \gamma(\alpha^g + x_0(\lambda + \delta\xi) - f) + f - rG(t) + (G_p + G_{pp}(t))(\frac{\alpha^g - \alpha^b + (x_0 - 1)\lambda}{\sigma})^2 \frac{1}{2},$$

the homogeneous solution of which can be found to be:

$$G(p) = \frac{\gamma(\alpha^g + x_0(\lambda + \delta\xi) - f) + f}{r} + K_1 e^{p\eta_1} + K_2 e^{p\eta_2}, \tag{25}$$

with $\eta_1 < 0 < \eta_2$. There are two relevant boundary conditions $G(\underline{p}) = 0$ and

$$\lim_{p \to \infty} G(p) = \frac{\gamma(\alpha^g + x_0(\lambda + \delta\xi) - f) + f}{r}, \tag{26}$$

the value that the manager would earn if her reputation is perfect and therefore she is never liquidated. The second boundary implies $K_2 = 0$, since $\eta_2 > 0$. K_1 is determined by:

$$\frac{\gamma(\alpha^g + x_0(\lambda + \delta\xi) - f) + f}{r} + K_1 e^{\underline{p}\eta_1} = 0,$$

then

$$G(p) = \begin{cases} \frac{\gamma(\alpha^g + x_0(\lambda+\delta\xi) - f) + f}{r}(1 - e^{(p-\underline{p})\eta_1}) & p > \underline{p}, \\ 0 & p \leq \underline{p}. \end{cases} \tag{27}$$

It follows that $\text{Sign}[G_{pp}] = \text{Sign}[G \times (-n_1^2)] < 0$. So managers of open-ended funds have concave value functions and are always averse to reputation risk. $\square$

Proposition 4. ***Manager's equilibrium portfolio choice.*** *Let $G \geq 0$, and $G_p \geq 0$. We let $x(p_t)$ denote the skilled manager's optimal portfolio choice, then $x(p_t)$ solves*

$$\max_x x(\lambda + \delta\xi) \tag{28}$$

$$s.t.\ x = \arg\max_{z \in [-1,1]} \left\{ \gamma z(\lambda + \delta\xi) + G_p \frac{\alpha^g + (x-1)\lambda - \alpha^b}{\sigma^2} z\lambda + \delta E\left[G\left(p + \frac{(x-1)\xi}{\omega^2}(\tilde{\omega} + z\xi - (x+1)\xi) \right) \right] \right\}, \tag{29}$$

where $E(\cdot)$ takes the expectation over the idiosyncratic shock $\tilde{\omega}$. The equilibrium choice can be broken into three cases:

(A) $x = -1$ if

$$\gamma(\lambda + \xi\delta) - 2(\frac{\xi}{\omega})^2 \delta E\left[G_p\left(p + \frac{2\xi}{\omega^2}(\xi + \tilde{\omega}) \right) \right] + G_p \frac{\alpha^g - 2\lambda - \alpha^b}{\sigma^2}\lambda \leq 0, \tag{30}$$

(B) if conditions in (A) does not hold, x is given by the lowest $z \in (-1, 1)$ that satisfies

$$z = 1 - \frac{\gamma(\lambda + \xi\delta) - G_p \frac{\alpha^g - \alpha^b}{\sigma^2}\lambda}{G_p \frac{\lambda^2}{2\sigma^2} + \frac{(\xi)^2}{\omega^2}\delta E\left[G_p\left(p + \frac{(z-1)\xi}{\omega^2}\left(\frac{(z-1)}{2}\xi + \tilde{\omega} \right) \right) \right]}, \tag{31}$$

(C) if problem in (B) has no solution then $x = 1$.

Proof. Given beliefs about the manager portfolio x, the right hand side of Equation (29) is the solution of the manager problem. Therefore this Equation (29) is a fixed point problem. In equilibrium, investors beliefs about the skilled manager portfolio have to be equal to the actual strategy implemented by the skilled manager. Because of strategic complementarities this fixed point problem might have multiple solutions. Intuitively, Equation (29) shows that a higher x increases the incentives for the manager to choose a higher z, so there is potential for multiplicity, especially when reputation concerns G_p are strong relative to compensation incentives $\gamma(\lambda + \delta\xi)$.

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 (29) as an Incentive Compatibility constraint. Because managers and investors can communicate, it is natural to assume that their coordination results in the mutually beneficial high-surplus equilibrium.

This proposition follows directly from three conditions that hold for the equilibrium: (1) the manager's first-order condition, (2) consistency of investors' beliefs about the manager's portfolio choice, and (3) return maximization.

The optimal choice divides into three regions depending on the intensity of reputational concerns. We first check if reputation concerns are weak enough so that $x(p_t) = -1$—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 x—because expected returns are decreasing in x—that satisfies the manager's first-order condition and it is consistency with investors beliefs. Case (C) is the corner solution where the manager chooses $x = 1$ since her incentives to take on the crash risk are too strong.

Intuition. The expression in case (A) in Proposition 4 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, than the manager's after-crash reputation will be very close to her 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 high after-crash reputation. Because long-term reputational concerns are decreasing in reputation ($G_{pp} < 0$), the marginal value of reputation after the crash is 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 her choices, they will be effective in keeping the manager from fully maximizing short-term performance. The long-term reputational concern is impacted not only by investors' beliefs x as in the case of short-term reputation concerns, but also directly by the actual manager's actual choice. Inspecting the distribution of crash reputation growth in the second line of Equation (29), we can see that as the portfolio increases towards 1 the manager expected reputation growth goes down. As a result, the manager expectation puts more weight on low reputation states happening after the crash, that is, states with higher reputation concerns $G_p \uparrow$. This increasing likelihood of a very bad performance and resulting reputation reduction disciplines the manager to not not go all the way to $x = 1$. □

Proposition 5. ***Unraveling of an open fund.*** *If the market for skill is perfectly competitive $\phi \to \infty$, the long-term trade is more profitable than the selection alpha $-(\lambda + \xi\delta) > \alpha^g$, and*

$$G_p(\underline{p}) \geq \frac{-\gamma(\lambda + \xi\delta)\sigma^2}{(\alpha^g - \alpha^b - 2\lambda)\lambda},$$

at liquidation, and there is perfect learning in the crash event $\omega \to 0$ (or no learning, $\omega \to \infty$), then *only a manager with reputation $p \to \infty$ $(P \to 1)$ can attract capital.*

Proof. The fact that the market is perfectly competitive implies that investors do not capture any upside from sticking with a manager. This means that the fund will be liquidated the first time the fund expected return dips negative. The condition $-(\lambda + \xi\delta) \geq \alpha^g$ implies that the skilled manager has non-positive expected return as $x \to 1$. Therefore, the manager portfolio cannot be expected to go to 1 as she approaches liquidation, because for any $p < \infty$ $(P < 1)$, the expected return for investors will be negative as they have to also consider the probability that the manager is unskilled. Therefore, the expected return of investing in the fund will always be negative at liquidation, what is inconsistent with the investors being indifferent. The condition on the slope of the value function G_p at liquidation guarantees that reputation concerns are strong enough at the liquidation threshold so that $x(\underline{p}) = 1$ is the only choice consistent with equilibrium. Therefore, both conditions together lead to the unraveling of an open fund, that is, only managers with perfect reputation are able to raise capital.

□

Proposition 6. ***Investors' liquidation policies.***

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

(a) The liquidation threshold $\underline{p}_l$ increases in the arrival rate of the outside investors, ϕ.

(b) The liquidation threshold $\underline{p}_l$ 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. The liquidation policy satisfies $V(\underline{p}, l_t = 0) = 1$. 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 imply smaller expected returns, 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. Since the liquidation decision is an option on the manager skill, this effect increases investor valuations. As $\phi \to \infty$ 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.

□

A.2 Hamilton-Jacobi-Bellman equations

Below we state the HJB equation of the manager in the case with an active market for skill and a lockup contract. As discussed in Sections 2.4 and 2.6 in this case the HJB becomes a system of two coupled ODEs. The ODEs describe the dynamics of the fund valuation in the active and inactive lockup states. In the Equation below, lines 1 and 2 are identical to the case of an open-ended fund in Equation (12). Lines 3 and 4 describe how the market for skill and the lockup provision regulates transitions between the two lockup states. Starting with line 3, it describes what happens when the fund is open and new investors arrive. If the manager reputation is below the reputation threshold required to attract locked capital, these new investors simply pay an access fee of $V(p,0)-1$ to replace the old investors in the fund, i.e. the manager remains managing an open fund. If when investors arrive, the manager reputation is high enough that $V(p,1) \geq 1$, i.e. $p_t \geq \underline{p}_1$, then she can attract locked capital and the fund lockup state transitions from inactive ($l=0$) to active ($l=1$) as new investors pay an access fee of $V(p,1)-1$ to replace the old investors in the fund but now with an active lockup. In the fourth line we see that if new investors arrive when the lockup is active, then again the new investors simply pay an access fee of $V(p,1)-1$ to replace the old investors in the fund and the manager remains managing an active lockup fund. The last term in line 4 describes the transition of the fund from the active lockup state to the inactive when the lockup expires with intensity $1/T$. The fifth line imposes the boundary condition that the manager ears zero if the fund is liquidated–what happens first time the reputation hits a value lower than $\underline{p}_0$ when the fund is open. The sixth lines describes the reputation dynamics when a crash hits the fund (see the discussion of Equation (12) for details).

$$
\begin{aligned}
rG^i(p_t,l_t) &= \max_{x_t^i}\Bigg\{ f + m(\alpha^i + x_t^i(\lambda+\delta\xi) - f) + G_p^i(p_t,l_t)\frac{E^I[\mathcal{D}(\alpha + x_t\lambda)]}{\sigma^2}\left(\alpha^i + x_t^i\lambda - E^I[\mathcal{A}(\alpha+x_t\lambda)]\right) \\
&\quad + \frac{1}{2}G_{pp}^i(p_t,l_t)\left(\frac{E^I[\mathcal{D}(\alpha+x_t\lambda)]}{\sigma}\right)^2 - \delta \mathrm{E}_t^i\left[G^i(p_{t+},l_t) - G^i(p_t,l_t)\right] \\
&\quad + \phi 1_{(l_t=0,\underline{p}_1\geq p_t)}\left(V(p_t,0)-1\right) + \phi 1_{(l_t=0,p_t\geq \underline{p}_1)}\left(G^i(p_t,1) - G^i(p_t,0) + V(p_t,1) - 1\right) \\
&\quad \phi 1_{(l_t=1,\geq p_t \geq \underline{p}_1)}\left(V(p_t,1)-1\right) + \frac{1}{T}1_{(l_t=1)}\left(G^i(p_t,0) - G^i(p_t,1)\right)\Bigg\}, \qquad (32) \\
G^i(p_t,0) &= 0,\ \forall p_t \leq \underline{p}_0 \qquad (33) \\
p_{t+} &= p_t + \frac{E^I[\mathcal{D}(x_t\zeta)]}{\omega^2}\left(x_t^i\zeta - E^I[\mathcal{A}(x_t\zeta)]\right) + \frac{E^I[\mathcal{D}(x_t\zeta)]}{\omega}y_{t+}^i, \qquad (34)
\end{aligned}
$$

We now describe the investors' HJB equation, which again is a system of coupled ODEs. The important difference between the investor and the manager HJB is the return dynamics and the transition across lockup states. Whereas the skilled manager knows her return 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. Furthermore, the transition across states is much simpler for an investor. The investors HJB can be seen below. The valuation effect of the transitions can be seen in line 3 of the Equation (35). They can always be pushed out of the fund and have a capital loss of $1 - V(p,l)$ if their manager have a reputation that is high enough to attract new capital what happens with intensity ϕ, i.e. they transition from invested to divested, and they can also transition from locked up to liquid with intensity 1/T when they have a capital gain of $V(p_t,0) - V(p_t,1)$. The left side of the first line is due to standard cash-flow discounting, the right side is the net cash flows received by the investors, i.e the net alpha of the fund before access fees paid at entry are factored in. The second line describes the valuation consequences of reputation movements. Because under the investors information set reputation is a martingale, there is no drift term, so the first term in the second line is due entirely due to the volatility, i.e. a convexity term due to the non-linearity in the investors value function. The second term is simply the reputation effect of a jump. The last three lines describe the boundary conditions of the ODEs.The first boundary condition pins down the value of investors holdings once she liquidates the fund, which is 1 the net asset value of fund assets. The second equation pins down the optimal time to liquidate the fund, and the last equation pins down the minimum reputation at which investors are happy to invest with a locked up fund.

$$\begin{aligned} rV(p_t, l_t) &= [r + (1-m)(E^I[\alpha^i + x_t^i(\lambda + \delta\xi)] - f)] \\ &+ \frac{1}{2}V_{pp}(p_t, l_t)\left(\frac{E^I[\mathcal{D}(\alpha + \Delta x_t \lambda)]}{\sigma}\right)^2 + \delta\left(E^I[V(p_{t+}, l_t)] - V(p_t, l_t)\right) \\ &+ \phi 1_{(p_t \geq \underline{p}_{l_t})}(1 - V(p_t, l_t)) + \frac{1}{T}1_{(l_t=1)}[V(p_t, 0) - V(p_t, 1)], \qquad (35) \\ V(p, 0) &= 1, \forall p \leq \underline{p}_0, \qquad (36) \\ V_p(\underline{p}_0, 1) &= 0, \qquad (37) \\ V(\underline{p}_1, 1) &= 1. \qquad (38) \\ & \qquad (39) \end{aligned}$$

A.3 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 x_t and the investors' investment policy for funds with and without lockups—are determined jointly. The state space consists of the manager's reputation in probability space ($P \in [0, 1]$) 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 $x(P, l) = -1$

and iterate to find what we denote the efficient solution $(G_{\text{ef}}, B_{\text{ef}}, V_{\text{ef}}, \underline{P}_0^{\text{ef}}, \underline{P}_1^{\text{ef}})$. The efficient liquidation threshold $\underline{P}_l^{\text{ef}}$ is the lower bound to the equilibrium liquidation threshold $\underline{P}_l$. Starting from the efficient policies we iterate on the HJB, but this time solving for the optimal portfolio $x(P, l)$ policy at each step.

The iteration procedure can be divided into steps:

1. Given choices x^{i-1}, solve for $\underline{P}_0^i, \underline{P}_1^i, V^i$ such that V^i satisfy the investors' HJB equation and the boundary conditions described in Section A.2,
2. Given $\underline{P}_0^i, \underline{P}_1^i, V^i$ and x^{i-1}, solve for G^i such that the manager's valuation satisfies the manager HJB equation described in Section A.2,
3. Given G^i, solve for x^i using Proposition 4,
4. If $\left|G^i - G^{i-1}\right| < \epsilon$ and $\left|V^i - V^{i-1}\right| < \epsilon$, 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 $x = -1$.