

Locked Up by a Lockup: Valuing Liquidity as a Real Option

Andrew Ang*

Columbia Business School

Nicolas P.B. Bollen

Vanderbilt University

Abstract

Hedge funds often impose lockups and notice periods to limit the ability of investors to withdraw capital. We model the investor's decision to withdraw capital as a real option and treat lockups and notice periods as exercise restrictions. Our methodology incorporates time-varying probabilities of hedge fund failure and optimal early exercise. We estimate a two-year lockup with a three-month notice period costs investors 1.5% of their initial investment. The magnitude is sensitive to a fund's age, expected return, and the liquidation cost upon failure. The cost of illiquidity can exceed 10% if the hedge fund manager suspends withdrawals.

JEL classification: G23, G13

Keywords: hedge fund lockup, withdrawal, redemption notice period, suspension clause

November 13, 2008

* Corresponding author: Andrew Ang, aa610@columbia.edu, Columbia Business School, 3022 Broadway, New York, NY 10027. Phone (212) 854-9154. Fax (212) 662-8474. The authors thank Greg van Inwegen, Hans Stoll, Jacob Sagi, and Neng Wang, as well as seminar participants at Cornell University, Vanderbilt University, and Virginia Tech, for helpful comments. Support from the Financial Markets Research Center is gratefully acknowledged.

Locked Up by a Lockup: Valuing Liquidity as a Real Option

1. Introduction

Hedge funds and funds-of-funds, along with many other alternative investment vehicles, place a variety of restrictions on the ability of investors to redeem their capital. A lockup requires an investor to wait a specified length of time after the initial deposit of capital, typically one to three years, before requesting a redemption. A notice period requires an investor to wait a specified length of time, typically one to three months, before a redemption request is processed. In addition, fund managers often have the authority to process only a fraction of a redemption request, known as a gate, or even to suspend redemptions altogether. As argued by Aragon (2007), the advantage of redemption restrictions is that they allow fund managers to invest in illiquid assets and earn an associated return premium. Redemption restrictions can levy an important cost, however, if they prevent investors from withdrawing capital before anticipated losses are realized.¹ Our goal is to develop a methodology to estimate the implied cost of redemption restrictions, thereby allowing investors to more accurately tabulate hedge fund fees.

We model the ability of an investor to withdraw capital as a real option. Upon exercise, the investor gives up ownership in the fund and receives a cash payoff per share equal to the fund's net asset value (hereafter "NAV"). The investor exercises the option when the investor's own valuation of a share of ownership in the fund falls below the NAV. We assume investors value the fund taking into account the probability of fund failure, liquidation costs, and the impact of future exercise decisions. Redemption

¹ The case of Amaranth Advisors is a good example, as described in "At Hedge Funds, Study Exit Guidelines," *Wall Street Journal* 10/23/06.

restrictions, such as lockups and notice periods, constrain the investor's ability to exercise, and their cost can be measured by the resulting reduction in value of the "liquidity option."²

Our approach has three key elements. First, for the liquidity option to have value there must be a difference between the NAV and an investor's valuation of ownership in the fund. The NAV is the value of the fund's portfolio reported by the fund manager on a given date. If the fund is invested in illiquid assets for which market prices are not readily available, the fund manager may employ subjective marking to model when computing the fund's NAV, and the NAV may be susceptible to managerial misreporting.³ We abstract from differences of opinion regarding the value of a fund's assets. Instead, in our model, an investor's valuation of ownership in the fund differs from the NAV because the latter reflects neither the capitalization of future managerial performance nor the probability of fund failure and the associated liquidation cost.

Second, we employ a data generating process (hereafter "DGP") for hedge fund returns that includes a normal regime with a constant expected return and an absorbing failure state in which investors are forced to accept a payout per share equal to a fraction of the fund's NAV. Motivated by Gregoriou (2002) and Grecu, Malkiel and Saha (2006), who find that most hedge funds stop reporting to databases because of failures following low returns, we use a log-logistic duration function to predict hedge fund failure, and allow hazard rates to depend on realized performance. Specifically, the probability of

² See Longstaff (1995), Scholes (2000), and Bollen, Smith, and Whaley (2004) for examples of liquidity options in other contexts.

³ See, for example, Asness, Liew and Krail (2001), Getmansky, Lo, Makarov (2004), and Bollen and Pool (2008).

fund failure, and hence the value of a liquidity option, changes over time as a function of fund age and performance, so that a fund with poor cumulative performance relative to its peers is more likely to fail.

Third, we model the payoff of investing in a hedge fund with and without a liquidity option using a binomial lattice that embeds time-varying probabilities of fund failure and naturally allows for early exercise. We measure the cost of a lockup as the difference between the value of a liquidity option that allows exercise at any time and another that does not permit exercise during the lockup. Similarly, we measure the cost of a notice period as the difference between the values of two liquidity options. The first liquidity option generates a payoff equal to the fund NAV immediately upon exercise. The second liquidity option generates an uncertain payoff because the redemption is not processed until the notice period has elapsed. During the notice period, the NAV can rise or fall and the fund may fail, hence the two liquidity options can have different payoffs. We also compute the combined cost of lockups and notice periods.

We estimate the cost of lockups and notice periods by calibrating our model to a large sample of hedge funds using the CISDM database. Parameters for the log-logistic duration function indicate that there is a 50% chance a fund will fail by age 78 months, though failures cluster in the first few years. Furthermore, a fund that falls one standard deviation below the cross-sectional mean has a 38% increased risk of failure. We estimate that the combined cost of a two-year lockup and a three-month notice period can exceed 1.5% of the initial investment, roughly the same as the average hedge fund management fee for one year. Furthermore, we show that a manager's discretion to block redemption requests using gate restrictions or suspension clauses generates an implied

cost of between 5% and 15% of the initial investment. This result indicates that hedge fund investors should be more concerned about the manager's authority to unilaterally void redemption procedures than standard lockups and notice periods.

The rest of the paper is organized as follows. Section 2 discusses the relations between our paper and existing research. Section 3 presents the DGP of hedge fund returns and shows how we value lockups and notice periods. The data and the calibrated DGP are described in Section 4. Section 5 computes the cost of lockups and notice periods over a range of inputs. Section 6 concludes.

2. Related literature

Our study of the ability of hedge fund investors to withdraw capital is related to the literature on closed-end mutual funds (hereafter "CEF") as well as existing work on hedge fund share restrictions.

Investors in our model exercise their redemption option when their own valuation of ownership in the fund falls below a hedge fund's NAV. Investors in CEFs can never redeem capital from a fund and instead trade shares of ownership on the secondary market, where they typically observe a difference between NAV and the market price. CEFs usually trade at a premium to NAV when first established, but spend most of their lives trading at a discount; hence the difference is labeled the CEF discount. Prior research has established several explanations for this phenomenon that are related to our model of hedge fund liquidity. Cherkes et al. (2008) model the CEF discount as the net of capitalized liquidity benefits and managerial fees. The liquidity benefits arise when the

fund is invested in illiquid securities that are costly to trade, so that an investor can reduce expenses by trading shares of the fund instead. Berk and Stanton (2007) model the CEF discount as the net of managerial ability to deliver abnormal returns and managerial fees. The typical pattern of a CEF trading at a premium initially, and then falling to a discount, can arise from a manager renegotiating fees after demonstrating ability. Like these papers, in our model there is a difference between the hedge fund NAV and the investor's valuation of the hedge fund, with the investor's valuation being the net of after-fee abnormal returns and the cost of fund failure.

Our paper differs fundamentally from the literature on CEF pricing for two reasons. First, unlike CEF investors, hedge fund investors do have the ability to exchange shares for NAV, although that ability is often restricted. Second, we model the investor's decision to redeem capital in the presence of time-varying probabilities of fund failure. Berk and Stanton (2007) do not consider fund failure at all. Cherkes et al. (2008) allow investors to force liquidation, which bears some resemblance to fund failure, although they assume it is never optimal to do so. In contrast, in our model an investor's decision to redeem is an optimal exercise of a real option which is affected by a fund failure process dependent on past performance. We do not model the strategic interactions in a liquidating event, as in Cherkes et al. (2008), but our empirically estimated failure rate's sensitivity to performance may capture some of this effect.

Our paper is also related to existing studies of the relation between redemption restrictions and hedge fund returns. Ding et al. (2007) show that redemption restrictions affect the empirical cross-sectional relation between aggregate capital flow and returns. Aragon (2007) documents that hedge funds with lockups have expected returns that are

4% – 7% per annum higher than hedge funds without lockups. Aragon interprets this difference as an illiquidity premium: lockups allow managers to invest in more illiquid securities and earn higher returns as a result. However, Aragon does not explicitly compute the cost of a lockup and cannot determine if the illiquidity premium is fair compensation. A related literature explores whether lockup provisions are a component of an optimal incentive contract for a fund manager, especially one investing in illiquid assets, as in Lerner and Schoar (2004).

Our paper is most closely related to Derman (2007), who models hedge fund returns using a three-state model in which hedge funds are good, sick, or dead. Derman, Park and Whitt (2007) extend this approach to allow for more complex Markov chain models. In both approaches, lockups prevent an investor from withdrawing capital from a sick fund and investing the proceeds in a good fund. Our valuation strategy differs from Derman (2007) in three important ways. First, Derman assumes investors swap capital invested in a poorly performing hedge fund for capital invested in a superior hedge fund, whereas we assume investors withdraw capital as cash. Thus, our approach explicitly models the actual decision that investors face. Second, we differentiate between lockups and notice periods, and develop a methodology that can estimate the cost of the two restrictions separately, or in combination. Third, we specify a failure state with a hazard rate that can depend on fund age and performance, as described next.

3. Methodology

In Section 3.1. we use a binomial lattice to model the dynamic evolution of fund NAVs conditional on a fund surviving. In Section 3.2. we augment the binomial lattice to

incorporate default probabilities. In Section 3.3. we explain how to use the augmented binomial lattice to value hedge funds with and without liquidity options, and how to estimate the cost of illiquidity by comparing hedge fund values when the liquidity option is restricted by a lockup, a notice period, or both.

3.1. Modeling fund NAVs

We assume that continuously compounded fund NAV returns are initially normally distributed and that this “normal regime” continues as long as the fund survives. We use a binomial lattice to model the evolution of a hedge fund’s NAV. Let $S_{t,j}$ denote the NAV, where t denotes the time step, running from 0 to T , and j denotes the level in the lattice, running from 1 to $t+1$ at time step t , with 1 being the highest, as depicted in Figure 1. We refer to the combination of time step t and level j as node (t, j) . The time between nodes is denoted by Δt . Date T represents the end of the hedge fund’s life if failure never occurs. This can be interpreted as the retirement of the hedge fund manager, or the feasible horizon of the fund’s investment strategy, and the purposeful unwinding of the hedge fund’s positions.

The geometry of the lattice is defined by the step size u and branch probability p , which are determined setting the mean and variance implied by the lattice equal to those of the hedge fund’s normal regime. With probability p the NAV increases from $S_{t,j}$ to $S_{t+1,j}$ where

$$S_{t+1,j} = S_{t,j}u, \tag{1}$$

with the multiplicative increase u and the probability p given by

$$\begin{aligned} u &= e^{\sigma\sqrt{\Delta t}} \\ p &= \frac{e^{\mu\Delta t} - u^{-1}}{u - u^{-1}}, \end{aligned} \tag{2}$$

where μ and σ are the mean and standard deviation of hedge fund returns in the normal regime. With probability $1 - p$ the NAV decreases from $S_{t,j}$ to $S_{t+1,j+1}$ where

$$S_{t+1,j+1} = S_{t,j}u^{-1}. \tag{3}$$

The parameters in (2) ensure that the distribution of NAV returns implied by the lattice converge to the assumed normal distribution as $\Delta t \rightarrow 0$.

The NAV is of central concern because this is the quantity received by the investor upon redemption. Furthermore, the NAV is empirically convenient because hedge funds report NAV returns to hedge fund databases. As described next, we model failure to be a function of fund age and performance, as measured by cumulative NAV returns relative to competing funds.

3.2. Failure process

Let D be the duration of a hedge fund, which we define as the random time that a fund fails, and at which point the manager liquidates the fund's remaining assets. Empirically, we measure duration as the time that a hedge fund manager stops reporting returns. While some hedge fund managers may stop reporting for good performance, the majority of funds cease reporting due to failure as argued by Ackermann, McEnally and Ravenscraft (1999) and Grecu, Malkiel and Saha (2006). In practice, a failing fund could

continue to survive for some period of time after the manager stops reporting, hence our measure likely understates durations. If the fund fails at node (t, j) , we assume that the fund NAV drops to a level $S_{t,j}l$ where l represents the proportion of pre-failure NAV that the manager is able to raise through liquidating asset sales, with $0 < l < 1$. Our model of the hedge fund failure process extends the Markov chain model employed by Derman, Park and Whitt (2007) to shift the baseline default intensity up and down by a performance covariate, as discussed below.

Denote the baseline density of durations as $f_b(t)$ with cumulative density function $F_b(t) = \int_0^t f_b(s) ds$. The baseline survival function $1 - F_b(t)$ is the unconditional probability of surviving up to at least time t , evaluated at time 0. The baseline hazard rate at time t is the probability of failure per time increment conditioned on surviving until t :

$$\lambda_b(t) = \lim_{\Delta t \rightarrow 0} \frac{\Pr(t \leq D \leq t + \Delta t | D \geq t)}{\Delta t} = \frac{f_b(t)}{1 - F_b(t)}. \quad (4)$$

We use a log-logistic hazard rate function to model hedge fund failures, following Grecu, Malkiel and Saha (2006), who find that the log-logistic distribution fits the empirical density of hedge fund durations better than other distributions. In Section 4, we present evidence that the log-logistic hazard rate function provides a tight fit to the empirical distribution of hedge fund failures in our sample. The log-logistic distribution is defined by two parameters λ and q with density, survival function, and hazard rate given by:

$$\begin{aligned}
f_b(t) &= \lambda q (\lambda t)^{q-1} / \left[1 + (\lambda t)^q \right]^2 \\
1 - F_b(t) &= 1 / \left[1 + (\lambda t)^q \right] \\
\lambda_b(t) &= \lambda q (\lambda t)^{q-1} / \left[1 + (\lambda t)^q \right].
\end{aligned} \tag{5}$$

We estimate λ and q by maximum likelihood. We separate hedge fund durations into n uncensored observations, for which the hedge fund manager stopped reporting prior to the end of the database, and m censored observations from “live” funds with observations through the end of the database. The likelihood of the data is then given by:

$$\prod_{i=1}^n f(t_i) \prod_{j=1}^m 1 - F_b(t_j) \tag{6}$$

with durations t_i for the n funds leaving the database and durations t_j for the m funds surviving until the end of the sample. This yields a log-likelihood of

$$n \ln(\lambda q) + (q-1) \sum_{i=1}^n \ln(\lambda t_i) - 2 \sum_{i=1}^n \ln\left(1 + (\lambda t_i)^q\right) - \sum_{j=1}^m \ln\left(1 + (\lambda t_j)^q\right). \tag{7}$$

We assume that failure rates also depend on hedge fund cumulative performance relative to other funds. Specifically, the hazard rate of an individual hedge fund equals the baseline hazard rate scaled up or down depending on the value of covariate z as follows:

$$\lambda(t; z) = \lambda_b(t) e^{z\beta}. \tag{8}$$

We follow standard practice and demean the covariate so that values above or below zero increase or decrease the hazard rate. We choose a performance-based covariate equal to the difference between the cumulative return of fund i at time t and the cross-sectional mean return. The difference is then scaled by the cross-sectional standard deviation. The

cross-sectional mean and standard deviation are computed using the cumulative returns of each fund when they are the same age as fund i at time t . The sensitivity of the failure rate to performance can be motivated in at least three ways. First and most importantly, Liang (2000), Brown, Goetzmann and Park (2001), and Jagannathan, Malakhov and Novikov (2006), among others, empirically document that liquidated hedge funds are more likely to be funds with poor past performance. Second, managers of funds with low cumulative returns are less likely to capture performance fees, since the NAV must recover to previously set high-water marks before the fees accrue. This provides managers of poorly performing funds a strong incentive to close those funds. Third, investors are more likely to withdraw capital from poorly performing funds, forcing the manager to liquidate assets, possibly leading to further reductions in NAV, again leading the manager to close those funds.⁴

The form of the proportional hazard rate in (8) is convenient because the coefficient β can be estimated independently from the baseline hazard rate as noted by Kalbfleisch and Prentice (2002). In particular, the relevant partial likelihood can be expressed as

$$\prod_{i=1}^n \left[e^{z_i \beta} \left[\sum_{k=1}^{N_i} e^{z_k \beta} \right]^{-1} \right], \quad (9)$$

where n is the number of uncensored observations, or fund failures, in the sample, z_i is the value of the performance covariate at failure of fund i , N_i is the number of funds in

⁴ As reported in Ding et al. (2007) academic evidence on the flow-performance relation in hedge funds is mixed, with prior research finding linear, convex, and concave relations between fund flow and performance. Ding et al. argue that the presence of redemption restrictions can explain these results.

the sample with durations at least as long as that of fund i , and z_k is the value of the performance covariate of fund k evaluated at age equal to that of fund i at the time of failure of fund i .

The hazard rate in (8) allows for the probability of failure to depend non-linearly on age and realized performance of the fund. In many option applications state-dependent payoffs lead to path dependence and cause the number of nodes in a lattice to explode. We avoid this by specifying the performance covariate so that it can be computed at each node in the lattice without knowledge of the path taken.

Let $\pi_{t,j}$ denote the probability of failure at node (t, j) , with the failure occurring prior to the return of the fund being realized between time t and $t+1$. At node (t, j) we numerically evaluate this by setting $\pi_{t,j} = \lambda_b(\text{age}_t + 0.5)\Delta t$ in the case of the base hazard rate and $\pi_{t,j} = \lambda_b(\text{age}_t + 0.5)e^{z_{t,j}\beta}\Delta t$ in the case of the proportional hazard rate, where Δt is the increment of time in the lattice, $z_{t,j}$ is the value of the performance covariate at node (t, j) , and age_t is the age of the fund at time t . The hazard rates over t to $t + \Delta t$ depend on the value of the covariate at t . Evaluating the baseline hazard rate at the midpoint amounts to taking the integral over t to $t + \Delta t$. Note that time $t=0$ corresponds to the investor's initial subscription to the fund rather than the fund's age. If the fund has already been in existence for some period of time, its cumulative return from inception of the fund impacts the value of $\pi_{0,1}$, since both the cumulative relative performance z as well as the baseline hazard function λ_b depend on the fund's age.

3.3. Valuing liquidity options

We have described the evolution of a hedge fund's NAV. But what is the fund worth to an investor? An investor's own valuation may differ from the NAV for a number of reasons. The investor's valuation incorporates the probability that the hedge fund will fail in the future and the associated liquidation cost.⁵ Furthermore, hedge fund investors generally have a real option to redeem their ownership and receive the NAV. The investor incorporates the probability of failure in her exercise decision: when the current NAV is greater than the expected, discounted hedge fund value next period, taking into account the probability of failure and the associated liquidation cost, then the investor should withdraw her capital.⁶

When an investor redeems, she is exchanging a share of ownership in the fund for its NAV, so one might think the redemption is the exercise of an exchange option, as in Margrabe (1978). Indeed, Derman (2007) explicitly models the exchange of an ownership in a bad fund for ownership in a good fund and computes the cost of a lockup as the inability to exchange a good fund for a bad fund during a fixed time period. In Margrabe's model, the option holder exchanges one risky asset for another, and both assets are governed by a distinct, but correlated, stochastic process. In our model, however, the investor's valuation of the fund is an explicit function of the NAV, so only the NAV's stochastic process is necessary. In addition, the investor exchanges ownership

⁵ Our pricing methodology does not assign any premium to an investor requiring immediate access to invested capital during the lockup period for exogenous reasons. Incorporating exogenous liquidity demands would increase the cost of lockups and notice periods.

⁶ If investors were identical they would all decide to withdraw simultaneously, forcing the manager to liquidate the fund so all investors would bear the liquidation cost. In practice, investors are not identical because some will be subject to lockups and others will not, depending on when they entered the fund. Some investors may also be better informed than other investors, or may possess different priors regarding the DGP of returns.

for cash, hence a more appropriate analog is a put option, in which the investor can sell the share of ownership back to the fund for the NAV. The NAV is of course constantly changing, so the liquidity option can be viewed as a put option with a variable exercise price.

We assume that investors are risk neutral and discount future payoffs using the risk-free rate. This can be motivated in two ways. First, many hedge funds advertise themselves as having absolute performance mandates and low correlation with standard asset classes, like the aggregate equity market. Indeed, Fung and Hsieh (1997) find that nearly half of hedge funds have R^2 of lower than 25% in Sharpe (1992) style regressions using major asset classes as regressands, and Bollen and Whaley (2008) report that the average hedge fund has adjusted R^2 less than 29% using a wide variety of style-based factors. Low levels of systematic risk correspond to low risk premia and a required rate of return approaching the risk-free rate. Second, investors investing in hedge funds must be wealthy.⁷ Carroll (2002) and Hurst and Lusardi (2004), among others, show that the very wealthy generally exhibit much lower levels of risk aversion and are substantially more willing to take risks than other households.

We expect that risk-averse investors would assign a higher cost to lockups and redemption periods than we compute. The literature on executive stock options (ESOs) shows that the value of an ESO to a manager who cannot short the underlying stock is affected by risk aversion.⁸ Higher risk aversion coincides with a higher likelihood that

⁷ Under Sections 3(c)(1) or 3(c)(7) of the Investment Company Act of 1940, investors in hedge funds must either be accredited investors, with net wealth greater than \$1 million, or qualified purchasers, with assets of over \$5 million. A typical required minimum investment in a hedge fund is \$5 or \$10 million.

⁸ See, among others, Kulatilaka and Marcus (1994), Detemple and Sundaresan (1999), and Murphy (1999).

constraints bind and this reduces the ESO value. Similarly, if higher risk aversion causes the restrictions on exercise imposed by lockups and notice periods to bind more often, then this would increase the cost of lockups and notice periods.

3.3.1. Hedge fund value with no liquidity option

Let $H_{t,j}$ denote the value of a hedge fund per share conditional on survival at node (t, j) from the perspective of a “passive” investor who does not possess a liquidity option. The passive investor is an artifice necessary to compute the value of liquidity options, as shown below. In addition, note that a fund manager can unilaterally impose restrictions on exercise that reduce the value of an investor’s liquidity option, and in the extreme can eliminate the liquidity option completely. Here is an excerpt from an actual partnership agreement that is typical of many suspension clauses employed by hedge funds:

The Fund may suspend redemptions and defer payment of redemption proceeds during any period in which disposal of all or part of the Fund's assets, or the determination of Net Asset Value, would not be reasonable or practical or would be prejudicial to the Fund or the Shareholders.

During times of high demand for the withdrawal of capital, for example, managers may temporarily suspend redemptions to avoid high transaction costs, such as price impact, that would be incurred when selling fund assets. Thus, $H_{t,j}$ represents the lower bound on the value of the hedge fund to an investor who possesses a liquidity option, a bound that is reached when the fund manager always suspends redemptions when it is optimal

for the investor to withdraw. In other words, we can interpret $H_{t,j}$ as the extreme case when a hedge fund manager disables an investor's liquidity option.

A passive investor will receive a payoff at the time of fund failure or at date T if the fund survives until the end of the fund's horizon. Assume that at the terminal set of nodes investors receive the NAV, i.e., $H_{T,j} = S_{T,j}$. This assumption could easily be relaxed to allow for the cost of unwinding existing positions in the fund at the terminal date. Prior to the terminal set of nodes, the value of the investment conditional on survival equals its expected, discounted value next period. Our assumption of risk-neutrality implies

$$H_{t,j} = \pi_{t,j} S_{t,j} l + e^{-r_f \Delta} (1 - \pi_{t,j}) (p H_{t+1,j} + (1-p) H_{t+1,j+1}), \quad (10)$$

where r_f is the risk-free rate. The first term on the RHS allows for the probability that the fund fails between t and $t+1$, which results in a loss due to the liquidation cost. The second term captures the value of the fund if failure does not occur. The recursion in (10) is repeated until the initial node in the lattice is reached, where the value of the hedge fund to the investor is $H_{0,1}$. Since the NAV reflects neither the probability of fund failure, nor the capitalization of fund returns in excess of the risk-free rate $H_{0,1} \neq S_{0,1}$.

3.3.2. Hedge fund value with unrestricted liquidity option

Let $O_{t,j}$ denote the value of the hedge fund per share at node (t, j) , again conditional on survival, but this time from the perspective of an "active" investor who possesses an unrestricted liquidity option. This liquidity option is defined as the ability to

exchange a share in the hedge fund for the NAV at any time prior to the terminal set of nodes. We estimate the cost of lockups and notice periods, which impose restrictions on exercise, from the basis of this unrestricted liquidity option. The unrestricted liquidity option will be exercised at a given node (t, j) if the expected, discounted hedge fund value next period is below the fund's prevailing NAV, $S_{t,j}$.

As before, we assume that the investor receives the NAV at the terminal set of nodes, that is $O_{T,j} = S_{T,j}$. Prior to the terminal set of nodes, the hedge fund value is the maximum of immediate exercise of the redemption option and the expected, discounted hedge fund value next period. As before, the expected, discounted hedge value next period incorporates the probability of failure between time t and $t+1$:

$$O_{t,j} = \max \left\{ S_{t,j}, \pi_{t,j} S_{t,j} l + e^{-r_j \Delta t} (1 - \pi_{t,j}) (p O_{t+1,j} + (1-p) O_{t+1,j+1}) \right\} \quad (11)$$

The initial value of the hedge fund to the investor with an unrestricted liquidity option is therefore $O_{0,1}$ and the initial value of the liquidity option is $O_{0,1} - H_{0,1}$.

3.3.3. Valuing a notice period

Let $O_{t,j}^{NP}$ denote the value of the hedge fund per share at node (t, j) , again conditional on survival of the fund, when the liquidity option is subject to a notice period. The hedge fund value is generally reduced by the notice period because when the option is exercised at node (t, j) the investor does not immediately receive a payoff, but rather must wait some period of time before receiving the prevailing NAV at that future date. The payoff may drop substantially if the fund fails while the investor waits, and indeed

this is relatively likely since the investor only chooses to exercise the option when the failure probability is high.

Consider a one-period delay between option exercise and receipt of the NAV. As before, we assume that the investor receives the NAV at the terminal set of nodes, i.e., $O_{T,j}^{NP} = S_{T,j}$. With a one-period notice, the option to redeem is irrelevant at the penultimate set of nodes, hence

$$O_{T-1,j}^{NP} = \pi_{T-1,j} S_{T-1,j} l + e^{-r_j \Delta t} (1 - \pi_{T-1,j}) (p O_{T,j}^{NP} + (1-p) O_{T,j+1}^{NP}). \quad (12)$$

Prior to the penultimate set of nodes, however, the hedge fund value is given by the maximum of the expected, discounted payoff of exercising the liquidity option and the expected, discounted hedge fund value next period. If the investor exercises at node (t, j) the payoff is the expected, discounted NAV at time $t+1$ since there is a one-period delay. We allow for the possibility that the fund fails while the investor is waiting, so that the expected, discounted NAV, denoted $ENAV$, is

$$ENAV_{t,j} = \pi_{t,j} S_{t,j} l + e^{-r_j \Delta t} (1 - \pi_{t,j}) (p S_{t+1,j} + (1-p) S_{t+1,j+1}). \quad (13)$$

Thus, the hedge fund value at node (t, j) is

$$O_{t,j}^{NP} = \max \left\{ ENAV_{t,j}, \pi_{t,j} S_{t,j} l + e^{-r_j \Delta t} (1 - \pi_{t,j}) (p O_{t+1,j}^{NP} + (1-p) O_{t+1,j+1}^{NP}) \right\}. \quad (14)$$

For no redemption notice, $ENAV_{t,j} = S_{t,j}$ and (14) reduces to (11).

Longer notice periods can be incorporated by replacing the expression in (13) with a sub-lattice that accounts for the possibility of failure at each time step. Consider an m -period delay in processing a redemption request that is made at node (t, j) . To

evaluate the hedge fund value with an m -period notice at node (t, j) , we compute the expected, discounted value of the fund taking into account the possibility the fund fails over the next m periods. The $ENAV$ with an m -period delay is computed in a recursive fashion, beginning with the set of $m+1$ NAVs that are possible in m periods assuming that the fund survives the notice period. In m periods, assuming the fund does not fail, the possible NAVs are $S_{t+m,j}$ through $S_{t+m,j+m}$. Moving back one time step, $ENAV$ is computed at each node as follows:

$$ENAV_{t+m-1,k} = \pi_{t+m-1,k} S_{t+m-1,k} l + e^{-r_f \Delta t} (1 - \pi_{t+m-1,k}) (p S_{t+m,k} + (1-p) S_{t+m,k+1}), \quad (15)$$

where k runs from j to $j+m-1$. Moving back one more time step, $ENAV$ is computed at each node in a similar fashion:

$$ENAV_{t+m-2,k} = \pi_{t+m-2,k} S_{t+m-2,k} l + e^{-r_f \Delta t} (1 - \pi_{t+m-2,k}) (p ENAV_{t+m-1,k} + (1-p) ENAV_{t+m-1,k+1}), \quad (16)$$

where k runs from j to $j+m-2$. The computation in (16) is then repeated for all remaining time steps in the notice period.

As before, we assume that the investor receives the NAV at the terminal set of nodes, that is $O_{T,j}^{NP} = S_{T,j}$. When the notice requires a wait of m periods, the option to file a notice does not exist in the m steps prior to the terminal date T . That is, the hedge fund value is computed as:

$$O_{t,j}^{NP} = \pi_{t,j} S_{t,j} l + e^{-r_f \Delta t} (1 - \pi_{t,j}) (p O_{t+1,j}^{NP} + (1-p) O_{t+1,j+1}^{NP}) \quad (17)$$

for $T-m \leq t \leq T-1$. Prior to this date, the investor decides whether or not to file a notice using the decision rule in (14).

The initial value of the hedge fund to the investor with a liquidity option restricted by a notice period is $O_{0,1}^{NP}$ and the initial value of the restricted liquidity option is $O_{0,1}^{NP} - H_{0,1}$. Recall the initial value of the liquidity option with no redemption restriction is $O_{0,1} - H_{0,1}$. Thus, the cost imposed on the investor by the notice period is the difference between the unrestricted and unrestricted option, i.e., $O_{0,1} - O_{0,1}^{NP}$. Hence, the cost of the redemption notice is the difference between two liquidity options. The first provides a payoff equal to the NAV immediately upon exercise. The second provides a payoff equal to the prevailing NAV at the end of the notice period, if the fund does not fail between the exercise of the option and the end of the notice period, or a payoff equal to a fraction of the prevailing NAV if the fund fails while the investor waits.

3.3.4. Valuing a lockup

A lockup prevents an investor from exercising her liquidity option prior to date L . Let $O_{t,j}^L$ denote the value of a hedge fund, at node (t, j) , with a liquidity option restricted by a lockup, and $O_{t,j}^{L,NP}$ denote the value subject to both a lockup and a notice period. After the lockup expires, there is no restriction, hence $O_{t,j}^L = O_{t,j}$ and $O_{t,j}^{L,NP} = O_{t,j}^{NP}$ for $t > L$, and hedge fund values can be computed as described in Section 3.3.3. For $t \leq L$, however, the investor cannot exercise the liquidity option, hence

$$\begin{aligned}
O_{t,j}^L &= \pi_{t,j} S_{t,j} l + e^{-r_j \Delta t} (1 - \pi_{t,j}) (p O_{t+1,j}^L + (1-p) O_{t+1,j+1}^L) \\
&\text{and} \\
O_{t,j}^{L,NP} &= \pi_{t,j} S_{t,j} l + e^{-r_j \Delta t} (1 - \pi_{t,j}) (p O_{t+1,j}^{L,NP} + (1-p) O_{t+1,j+1}^{L,NP})
\end{aligned} \tag{18}$$

for $t \leq L$. The lockup restricts exercise of the liquidity option in the same way that a vesting period prevents exercise of ESOs, as in Hull and White (2004).

The initial value of a hedge fund with a liquidity option subject to a lockup is therefore $O_{0,1}^L$ and the initial value of the restricted liquidity option is $O_{0,1}^L - H_{0,1}$. The ex-ante cost of the lockup itself is the difference between the unrestricted and restricted options, i.e., $O_{0,1} - O_{0,1}^L$. Similarly, the initial value of a hedge fund with a liquidity option subject to both a lockup and a notice period is $O_{0,1}^{L, NP}$ and the initial value of the restricted liquidity option is $O_{0,1}^{L, NP} - H_{0,1}$. The combined ex-ante cost of the lockup and notice period is the difference between the unrestricted and restricted options, i.e., $O_{0,1} - O_{0,1}^{L, NP}$.

3.3.5. Summary

Table 1 contains a summary of notation describing model parameters, including expressions for the value of liquidity options and the cost of notice periods and lockups. The value of a liquidity option is computed as the difference between hedge fund value with the liquidity option and hedge fund value from the perspective of a passive investor with no liquidity option. The cost of a notice period or a lockup is computed as the difference between the hedge fund value with an unrestricted liquidity option and the hedge fund value with a liquidity option subject to the restriction. As mentioned in Section 1, another type of exercise restriction that managers can impose is a gate, which limits the quantity of redemption requests that are accommodated in any period. Gate restrictions are typically invoked when there are an unexpected large number of redemptions requested by investors and the fund restricts, usually on a pro-rata basis, the

amount of money each investor can receive. Our analysis can be easily modified to include the effect of a gate, but we choose to focus on lockups and notice periods since the implementation of a gate is discretionary and may involve strategic games among investors to redeem from a fund with a high probability of failure before other investors redeem.

4. Data and Parameter Estimates

The hedge fund data used in our empirical analysis are from the Center for International Securities and Derivatives Markets (CISDM) database. The sample period runs through December 2005. The CISDM database includes live and defunct hedge funds, funds of funds, CTAs, commodity pool operators, and indices. We eliminate indices, since they have no partners and hence no lockup feature. There are 4,260 defunct funds and 4,272 live funds in the sample, with a total of 504,979 monthly observations.

Figure 2 shows the percentage of defunct funds at each possible duration, as well as the percentage of live funds at each possible history length. The most common durations for defunct funds lie in the two to four year range, suggesting that many funds fail early in their lives. The most common history length for live funds is one year or less, a result of the tremendous growth in the industry. Table 2 lists the interquartile range of the history lengths of live funds, defunct funds, and the full sample. The ranges are similar with medians of 44 months for live funds and 47 months for defunct funds. The hedge fund data are right-tailed censored because we know only the history of returns up until a fund stops reporting, or until the end of the database is reached. Hence the

expected duration of a fund will likely be much longer than these medians, and this is verified by maximum likelihood estimates of the hazard rate function as described below.

Cross-sectional averages of annualized summary statistics of returns are listed in Table 3. We require at least 24 observations for a fund to be included – a total of 3,287 defunct funds and 3,023 live funds, covering 476,178 monthly returns, are represented. The performance of defunct funds is clearly inferior, as is to be expected if poor performance is a predictor of fund failure. For example, defunct hedge funds have an average return of 11.86% compared to 13.87% for live hedge funds. Sharpe ratios of defunct hedge funds average 0.63 versus 1.27 for live hedge funds. Note there is also substantial variation across fund types. In Panel C, for example, hedge funds have volatility of 15.74%, whereas CTAs have volatility of 22.57%. In unreported analysis, wide variation also exists across subsets of hedge funds formed by strategy. Thus, when computing the value of liquidity options, and the cost of lockups and notice periods, it will be important to consider a wide range of parameters since there is large heterogeneity across hedge funds. As a base case, we use expected return of 10% and volatility of 15% for the “normal regime.”

Table 4 lists parameter estimates of the hazard rate function. Parameters of the baseline log-logistic function in (5) are estimated very precisely, with $\lambda = 0.0129$ and $q = 1.6517$. The duration at which the unconditional probability of survival is 50% is given by λ^{-1} , which equals approximately 78 months. Thus, taking into account the large number of censored observations in the sample, the expected duration of a fund is indeed much longer than the median 44-month duration of live funds reported in Table 2. Figure 3 compares the actual number of hedge funds with uncensored duration t to the predicted

number based on the parameter estimates of λ and q . The predicted number equals the hazard rate evaluated at duration t times the total number of funds, both live and defunct, with history length at least t . The fit is good, with the sharp peak for short durations consistent with the empirical distribution of durations of defunct funds in Panel A of Figure 2. This result suggests that fund age is a significant determinant of the probability that a manager will stop reporting to the database. The other determinant in our specification is the relative cumulative fund performance which shifts the baseline hazard rate up or down. Table 4 reports the maximum likelihood estimate of β is -0.3237 with standard error 0.0059. Thus, a cumulative return that is one standard deviation below the mean increases the hazard rate by about 38%.

Upon failure, we assume as a base case that investors receive a payoff of $l = 50\%$ of the prevailing NAV of the fund, reflecting additional loss of asset value during liquidation. The 50% liquidation cost is based on results reported in Ramadorai (2008), who analyzes a sample of transactions on a secondary market for hedge fund investments conducted on Hedgebay. During 66 “disaster” transactions, involving fraud or collapse, the average discount of transaction price to NAV is 49.6%.

5. Value of liquidity options and the cost of restrictions

We estimate hedge fund values with and without liquidity options, and compute costs of redemption restrictions, over a wide range of parameter values. Section 5.1. derives a threshold failure probability at which an investor is indifferent about exercising

the redemption option. The threshold failure probability provides insights regarding the sensitivity of our estimates to different parameter values. Section 5.2. presents the results.

5.1. Effect of fund failure

The cost of lockups and notice periods is generated by their restrictions on the investor's exercise decision, and the resulting decrease in the value of the investor's liquidity option. To gain some insight into how parameters of the model affect the cost of restrictions, we determine when the investor is indifferent about exercising the redemption option. From the recursion in (11), if the following relation holds

$$S_{T-1,j} = \pi_{t,j} S_{T-1,j} l + e^{-r_f \Delta t} (1 - \pi_{t,j}) (p S_{T-1,j} \mu + (1 - p) S_{T-1,j} d) \quad (19)$$

then the investor is indifferent between holding the fund and exercising at time $T-1$. Equation (19) reveals the fundamental determinant of liquidity option value: at node (t, j) the investor receives with probability $\pi_{t,j}$ a gross return equal to the liquidation payoff upon default, l , and with probability $1 - \pi_{t,j}$ the investor receives an expected return equal to that of the normal regime, μ . The investor is indifferent about exercising when these two outcomes offset each other.

For a given liquidation payoff l and expected return μ , there exists an analytic steady-state threshold value for failure probability at which an investor is indifferent about exercising. Substituting for p and d in (19) yields the critical failure probability

$$\pi_{T-1,j}^* = \frac{e^{(\mu - r_f) \Delta t} - 1}{e^{(\mu - r_f) \Delta t} - l}. \quad (20)$$

The expression in (20) shows that the threshold failure probability is increasing in both expected return and the proportion of assets retained in case of failure. For example, with a liquidation payoff l equal to 50%, an annual risk-free rate of 4%, and an expected return of 12%, the threshold failure probability is 1.3%. We can determine whether it is optimal to exercise by comparing a fund's actual failure probability to this threshold.

In Figure 4, we plot the probability of failure at each time step in a 120-month lattice in the absence of a performance covariate. The probability is given by the baseline hazard rate integrated over the times between nodes, which is approximated by the hazard rate evaluated at the time between nodes:

$$\pi(t) = \lambda q (\lambda t)^{q-1} / \left[1 + (\lambda t)^q \right] \Delta t \quad (21)$$

where the parameters $\lambda = 0.0129$ and $q = 1.6517$ are from Table 4 and Δt equals one month. Note that the hazard rate peaks at about 1.1% at a horizon of 60 months and slowly decays thereafter, always below the 1.3% threshold computed above. Hence, for these parameters, and in the absence of a performance covariate, it is never optimal to exercise.

In our model, however, the default intensity is affected by the performance covariate; hence exercise can be optimal for all funds depending on the realized returns of the fund. Figure 5 shows the probability of failure at each node in a 36-month lattice where the parameters of the hazard rate and performance covariate are set to the estimates listed in Table 4. Figure 5 shows the failure probability for a fund with expected return of 10% and three levels of return volatility, corresponding roughly to the cross-sectional average volatility (15%) as well as the 75th and 95th percentiles (25% and 40%,

respectively). Holding all else equal, the higher the fund volatility, the greater is the failure probability following poor performance. Intuitively, with greater volatility, the fund can be further away from the cumulative cross-sectional mean and the effect of the performance covariate becomes larger.⁹ Note that in many regions of all three lattices, the failure probability is many times higher than threshold failure probability described in (20), indicating that exercise is often optimal, and hence imposing restrictions may be costly.

In summary, the value of a liquidity option depends on the likelihood that a fund's failure probability will exceed the threshold. We have shown that failure probabilities are highest for relatively young funds; hence the liquidity option value and the cost of restrictions will likely peak for young funds as well. From (20), the threshold failure probability is increasing in both μ and l ; hence the liquidity option will be more valuable with lower expected returns and a greater loss upon liquidation. We now verify these predictions.

5.2. Value of liquidity options

Table 5 lists the combined costs of lockups and notice periods when initial fund NAV is \$100, returns are normally distributed with annual expected return of 10% and volatility of 15%, the fund has a ten-year horizon, and fund failures incur a 50% loss upon liquidation. Panel A shows results for new funds, whereas Panel B shows results

⁹ Consistent with the assumption of risk neutrality, if the hazard rate were purely a deterministic function of time, then changing the volatility would not change the value of the hedge fund or liquidity options. However, the impact of the performance covariate is increasing in volatility, hence the value of the hedge fund and liquidity options are affected by the level of volatility. In unreported analysis, we find that the impact on liquidity options and the cost of restrictions is relatively minor.

when funds have age equal to 24 months at $t = 0$. Panel A shows that costs are increasing in the length of the lockup and notice period. For a notice period of three months, the combined cost of restrictions ranges from \$0.09 for a one-year lockup to \$2.42 for a five-year lockup. For a lockup of three years, the combined cost of restrictions ranges from \$0.59 for a one-month notice period to \$0.89 for a five-month notice period. These results indicate that the length of the lockup has much more impact on the combined cost of restrictions than the length of the notice period. Panel B shows that slightly older funds generally have higher restriction costs – the intuition for this is that failure probabilities peak for funds that are a few years old, as indicated by Figure 2. For a lockup of three years, for example, the combined costs range from \$1.18 to \$1.44 over the notice periods, roughly double the cost when funds are new. Interestingly, for five-year lockups, the combined costs are slightly lower in Panel B. Recall that the cost of restrictions is computed as the difference between hedge fund values with unrestricted and restricted liquidity options, and both of these can decrease as failure probabilities change.

To illustrate this, Table 6 reports three panels of hedge fund values corresponding to hedge funds with an unrestricted liquidity option, hedge funds with a liquidity option restricted by a two-year lockup and a three-month notice period, and hedge funds with no liquidity option. Values are shown over a range of fund ages and horizons, T . Panel A shows that, for a given horizon, fund values display a U-shaped pattern with respect to fund age. Since failure probabilities peak for funds aged two to three years, hedge fund value initially drops with fund age. However, conditional on a fund surviving, failure probabilities subsequently fall and fund values increase. At a ten-year horizon, for example, the fund value starts at \$106.04, drops to \$100.80 at an age of 48 months, and

then increases to \$107.09 for an age of 96 months. The unconditional expected return of the fund, taking into account the return in the normal regime as well as the probability and cost of failure, is highest for the oldest funds. The pattern is present in Panels B and C as well.

A comparison of Panels A and B shows the decrease in fund value when restrictions are imposed, i.e. the cost of restrictions. For example, at an age of 24 months and a ten-year horizon, the cost is \$101.06 less \$100.35 or \$0.71, equal to the cost listed for a two-year lockup and a three-month notice period in Panel B of Table 5. Note that this cost shrinks to \$0.53 for fund age of 48 months in Table 6. The reason is that both hedge fund values drop with the increased risk of failure, but the unrestricted liquidity option drops more. Note also in Table 6 that fund values are generally increasing in horizon. The intuition is that, in most cases, the unconditional expected return of the fund, taking into account the expected return in the normal regime as well as the probability and cost of failure, is greater than the risk-free rate, so that the greater the horizon the greater the present value of fund payoffs. For young funds, the impact of T is relatively minor. For funds aged 72 or 96 months, though, the longer horizons have much higher values. Funds with an unrestricted liquidity option, for example, have values ranging from \$102.86 to \$111.98 as horizon increases from 72 months to 168 months.

We display the value of liquidity options and the cost of restrictions as a function of expected return, hazard rate, and loss upon liquidation, in Figure 6. Panel A shows results over a range of expected returns and a range of hazard rates, in which the empirical estimate of the parameter λ is scaled by the factor listed in the charts. The liquidity option value is decreasing in expected return, since the failure probability is

reduced via the performance covariate. Liquidity option value is increasing in λ because the baseline hazard rate increases. Note that the cost of restrictions is essentially zero at all levels of expected return when the λ is low. The reason for this is that at very low hazard rates, the probability that exercise will be optimal during the lockup period is negligible. Panel B shows results when l , the proportion of assets retained upon fund failure, is varied. Liquidity option value and the cost of restrictions are both decreasing in l since the benefit of exercise is lower when losses due to failure are lower.

Table 7 reports hedge fund values with and without restrictions, over a range of expected returns in the normal regime. Panel A lists hedge fund values when the fund has age = 0 at time $t = 0$. The passive value of the hedge fund with no liquidity option, $H_{0,1}$, ranges from \$72.59 when $\mu = 6\%$ to \$146.13 when $\mu = 14\%$. Discounts to NAV at lower levels of expected return reflect the probability of failure and the subsequent loss upon liquidation. Premiums to NAV at higher levels of expected return occur because the risk-neutral investor values the hedge fund by discounting at the risk-free rate. For an investor with an unrestricted liquidity option, $O_{0,1}$, the hedge fund value has no discount to NAV since the investor only remains invested when the expected holding period return (taking into account the expected return in the normal regime and the probability of fund failure) exceeds the risk-free rate. As the expected return in the normal regime increases, the value of the hedge fund increases as well. At $\mu = 10\%$, for example, $H_{0,1} = \$101.32$ and $O_{0,1} = \$106.04$, making the unrestricted liquidity option worth $O_{0,1} - H_{0,1} = \$4.72$.

Introducing notice periods has only a small effect on hedge fund value because the probability of failure during the notice period is small. Similarly, when $\mu > 10\%$, the

lockup restriction has little effect on the hedge fund value relative to $O_{0,1}$ because at this level of expected return the probability that exercise is optimal during the lockup period is extremely low. This is also true for the value of the hedge fund when the investor possesses a liquidity option subject to both lockup and notice period restrictions, $O_{0,1}^{L,NP}$. For example, when $\mu = 10\%$, the restricted liquidity option with a lockup and notice period is $O_{0,1}^{L,NP} = \$105.96$ whereas, for an unrestricted liquidity option, $O_{0,1} = \$106.04$, making the cost of the lockup and notice period small, at $O_{0,1} - O_{0,1}^{L,NP} = 0.08$. In contrast, the restriction imposed by a lockup on the hedge fund value causes a noticeable drop from the unrestricted liquidity option value for low expected returns because the probability that exercise is optimal during the lockup period is much higher. For $\mu = 8\%$, the value of the fund with a lockup and a notice period is $O_{0,1}^{L,NP} = \$100.26$, just above the initial NAV of \$100, so that is just worth investing in the fund. With the restrictions removed, $O_{0,1} = \$101.71$, hence the restrictions cost \$1.45 or about 1.5% of the initial investment.

Panel B lists values when the fund is 24 months old at time $t = 0$ and has cumulative return equal to the cross-sectional mean at that time, so that the performance covariate has no impact on failure probability initially. In all cases, the values are lower than in Panel A because the initial failure probability is higher. Panel B shows that investors with an unrestricted liquidity option would in fact optimally exercise immediately, so that the value of the fund equals the NAV, for all levels of expected return less than 10%.

Recall that in addition to lockups and notice periods, limitations on the quantity of a redemption processed, known as gates, and, in some cases, outright suspensions of redemption ability can restrict an investor's real option to withdraw capital.^{10,11} In the limit, gates and suspensions ultimately eliminate the investor's liquidity option. An investor who believes they are investing at time $t = 0$ in a fund with a lockup and notice period with value $O_{0,1}^{L,NP}$ may in fact be buying a fund with no liquidity option at all, which is worth $H_{0,1}$. Panel A of Table 7 shows that for relatively low expected returns, such as $\mu = 8\%$, the potential cost of redemption suspension can be enormous, with $H_{0,1} = \$85.37$ and $O_{0,1}^{L,NP} = \$100.26$. This implies that the investor is receiving an asset worth about 15% less than NAV, rather than one worth about par when the liquidity option is honored.

6. Conclusion

We model the investor's decision to redeem capital from a hedge fund as a real option, and develop a methodology to value the cost of lockups and notice periods. An investor who is always able to redeem from the fund and receive the prevailing NAV has an unrestricted liquidity option. Lockups and notice periods are exercise restrictions that reduce the value of the liquidity option. The cost of the restrictions is estimated by the resulting reduction in the value of the liquidity option investors possess.

¹⁰ See, for example, "Ore Hill Closes Fund to Client Withdrawals," *Wall Street Journal* 8/23/08.

¹¹ Gates could be accommodated in our lattice by modifying the payoffs that occur when redemptions are processed after a notice period has elapsed – instead of the payoff occurring all at once, they occur over a sequence of nodes following the gate. Similarly, suspensions of redemption ability could be represented by extending the lockup beyond the stated horizon.

We value the liquidity options using a lattice that accounts for the possibility of early exercise and incorporates time-varying probabilities of fund failure that vary with fund age and fund performance. We find that typical parameter values can generate costs of 1.5% of initial NAV for a two-year lockup and a three-month notice period. These estimates are well below the liquidity premium that hedge fund investors gain, as reported by Aragon (2007), so hedge funds are able to earn returns that more than compensate investors for the cost of lockups and notice periods. We show that the cost of restrictions is sensitive to fund-specific attributes such as age, expected return, and the loss generated by liquidation of fund assets. Furthermore, our estimates are based on failure probabilities that use relatively little fund-specific information. Funds with higher failure rates will have more valuable liquidity options, and hence restrictions thereon will be more costly. We leave for future research the exercise of estimating the cost of restrictions for each fund in our sample, as well as an empirical study of the relation between strategy, the presence of restrictions, and their cost.

When fund managers can unilaterally suspend an investor's real option to redeem, we show that the cost of illiquidity can be as much as 15% of initial fund NAV. This result suggests that hedge fund investors should be more concerned about the discretion asserted by fund managers in their partnership agreement, and conditions under which redemption suspensions can be imposed, rather than by the standard terms of lockup and notice periods.

References

- Ackermann, C., McEnally, R., Ravenscraft, D., 1999. The performance of hedge funds: risk, return, and incentives. *Journal of Finance* 54, 833-874.
- Amihud, Y., Mendelson, H., 1986. Asset pricing and the bid-ask spread. *Journal of Financial Economics* 17, 223-249.
- Aragon, G. O., 2007. Share restrictions and asset pricing: evidence from the hedge fund industry. *Journal of Financial Economics* 83, 33-58.
- Asness, C., Krail, R., Liew, J., 2001. Do hedge funds hedge? *Journal of Portfolio Management* 28, 6-19.
- Berk, J. B., Stanton, R. H., 2007. Managerial ability, compensation, and the closed-end fund discount. *Journal of Finance* 62, 529-556.
- Bollen, N. P. B., Pool, V., 2008. Conditional return smoothing in the hedge fund industry. *Journal of Financial and Quantitative Analysis* 43, 267-298.
- Bollen, N. P. B., Whaley, R. E., 2008. Hedge fund risk dynamics: implications for performance appraisal. *Journal of Finance*, forthcoming.
- Bollen, N. P. B., Smith, T., Whaley, R. E., 2004. Modeling the bid/ask spread: measuring the inventory-holding premium. *Journal of Financial Economics* 72, 97-41.
- Brown, S. J., Goetzmann, W. N., Park, J., 2001. Careers and survival: competition and risk in the hedge fund and CTA industry. *Journal of Finance* 56, 1869-1886.
- Carroll, C. D., 2002. Portfolios of the rich. In *Household Portfolios*, Guiso, L., M. Haliassos, and T. Jappelli, eds., MIT Press, Cambridge MA.
- Cherkes, M., Sagi, J., Stanton, R., 2008. A liquidity-based theory of closed-end funds. *Review of Financial Studies*, forthcoming.
- Cox, J. C., Ross, S. A., Rubinstein, M., 1979. Option pricing: a simplified approach. *Journal of Financial Economics* 7, 229-263.
- Derman, E., 2007. A simple model for the expected premium for hedge fund lockups. Working paper, Columbia University.
- Derman, E., Park, K. S., Whitt, W., 2007. Markov chain models to estimate the premium for extended hedge fund lockups. Working paper, Columbia University.
- Detemple, J., Sundaresan, S., 1999. Nontraded asset valuation with portfolio constraints. *Review of Financial Studies* 12, 835-872.
- Ding, B., Getmansky, M., Liang, B., Wermers, R., 2007. Investor flows and share restrictions in the hedge fund industry. Working paper, University of Massachusetts.
- Fung, W., Hsieh, D. A., 1997. Empirical characteristics of dynamic trading strategies: the case of hedge funds. *Review of Financial Studies* 10, 275-302.
- Fung, W., Hsieh, D. A., 2001. The risk in hedge fund strategies: theory and evidence from trend followers. *Review of Financial Studies* 14, 313-341.

- Getmansky, M., Lo, A. W., Makarov, I., 2004. An econometric model of serial correlation and illiquidity in hedge fund returns. *Journal of Financial Economics* 74, 29-609.
- Greco, A., Malkiel, B. G., Saha, A., 2006. Why do hedge funds stop reporting their performance? CEPS working paper 124.
- Gregoriou, G. N., 2002. Hedge fund survival lifetimes. *Journal of Asset Management* 3, 237-52.
- Hull, J., White, A., 2004. How to value employee stock options. *Financial Analysts Journal* 60, 114-119.
- Hurst, E., Lusardi, A., 2004. Liquidity constraints, household wealth, and entrepreneurship. *Journal of Political Economy*, 112, 319-347.
- Jagannathan, R., Malakhov, A., Novikov, D., 2006. Do hot hands persist among hedge fund managers? An empirical evaluation. NBER working paper 12015.
- Kahl, M., Liu, J., Longstaff, F. A., 2003. Paper millionaires: how valuable is stock to a stockholder who is restricted from selling it? *Journal of Financial Economics* 67, 385-410.
- Kalbfleisch, J. D., Prentice, R. L., 2002. *The Statistical Analysis of Failure Time Data*, 2nd edition. John Wiley & Sons, Hoboken, NJ.
- Kulatilaka, N., Marcus, A., 1994. Valuing employee stock options. *Financial Analysts Journal*, 46-56.
- Lerner, J., Schoar, A., 2004. The illiquidity puzzle: theory and evidence from private equity. *Journal of Financial Economics* 72, 3-40.
- Liang, B., 2000. Hedge funds: the living and the dead. *Journal of Financial and Quantitative Analysis* 35, 309-326.
- Longstaff, F. A., 1995. How much can marketability affect security values? *Journal of Finance* 50, 1767-1774.
- Margrabe, W., 1978. The value of an option to exchange one asset for another. *Journal of Finance* 33, 177-186.
- Merton, R. C., 1980. On estimating the expected return on the market: an exploratory investigation. *Journal of Financial Economics* 8, 323-361.
- Murphy, K., 1999. Executive compensation. In: Ashenfelter, O., Card, D. (Eds.), *Handbook of Labor Economics*. North-Holland.
- Ramadorai, T., 2008. The secondary market for hedge funds and the closed-hedge fund premium. University of Oxford working paper.
- Scholes, M. S., 2000. The near crash of 1998: crisis and risk management. *American Economic Association Papers and Proceedings* 90, 17-21.
- Sharpe, W. F., 1992. Asset allocation: management style and performance management. *Journal of Portfolio Management* 18, 7-19.

Table 1. Nomenclature

Listed are the symbols used to denote hedge fund value with and without a liquidity option, the value of liquidity options, and the cost of redemption restrictions.

H	=	Fund value with no liquidity option
O	=	Fund value with unrestricted liquidity option
O^L	=	Fund value subject to lockup
O^{NP}	=	Fund value subject to notice period
$O^{L,NP}$	=	Fund value subject to lockup and notice period
$O - H$	=	Value of unrestricted liquidity option
$O^L - H$	=	Value of liquidity option subject to lockup
$O^{NP} - H$	=	Value of liquidity option subject to notice period
$O^{L,NP} - H$	=	Value of liquidity option subject to lockup and notice period
$O - O^L$	=	Cost of lockup
$O - O^{NP}$	=	Cost of notice period
$O - O^{L,NP}$	=	Combined cost of lockup and notice period

Table 2. Durations.

Listed is the interquartile range of durations in months of hedge funds in the 2005 CISDM hedge fund database.

	No. of Funds	25 th	50 th	75 th
Live	4,272	20	44	84
Defunct	4,260	25	47	80
All	8,532	23	45	82

Table 3. Summary statistics.

Listed are annualized summary statistics of monthly returns of funds in the 2005 CISDM database. Only funds with at least 24 observations are included. Fund types are hedge funds, HF, funds of funds, FOF, commodity trading advisors, CTA, and commodity pool operators, CPO. Summary statistics include mean, μ , standard deviation, σ , Sharpe ratio, SR , skewness, $Skew$, and excess kurtosis, $Kurt$.

Type	No. of Funds	μ	σ	SR	$Skew$	$Kurt$
HF	1,685	11.86%	18.59%	0.63	-0.02	4.14
FOF	388	7.22%	9.65%	0.58	-0.31	4.13
CTA	548	12.83%	23.64%	0.31	0.64	3.87
CPO	666	7.29%	19.46%	0.15	0.34	3.60
All	3,287	10.55%	18.56%	0.47	0.13	3.99

Type	No. of Funds	μ	σ	SR	$Skew$	$Kurt$
HF	1,450	13.87%	12.43%	1.27	0.15	3.55
FOF	1,025	8.54%	5.91%	1.26	-0.24	2.50
CTA	303	15.11%	20.62%	0.60	0.54	2.88
CPO	245	10.76%	18.54%	0.47	0.46	2.01
All	3,023	11.93%	11.53%	1.14	0.08	3.00

Type	No. of Funds	μ	σ	SR	$Skew$	$Kurt$
HF	3,135	12.79%	15.74%	0.92	0.06	3.87
FOF	1,413	8.18%	6.94%	1.08	-0.26	2.94
CTA	851	13.64%	22.57%	0.41	0.60	3.52
CPO	911	8.22%	19.21%	0.23	0.37	3.17
All	6,310	11.21%	15.19%	0.79	0.11	3.51

Table 4. Duration parameters.

Listed are parameter estimates for the following hazard rate to model the probability of hedge fund failure:

$$\lambda(t; z) = \lambda q (\lambda t)^{q-1} / [1 + (\lambda t)^q] \exp(z\beta)$$

where λ , q , and β are parameters, t is the age of the fund, and z is the value of a performance score which equals the number of cross-sectional standard deviations the fund's cumulative return is from the cross-sectional mean. Parameters are estimated by maximum likelihood using the 2005 CISDM hedge fund database.

	λ	q	β
Estimate	0.0129	1.6517	-0.3237
Std Error	0.0002	0.0200	0.0059

Table 5. Combined Cost of Lockup and Notice Period.

Listed are the combined costs of lockups and notice periods of different lengths, in months, per share of a hedge fund with initial NAV of \$100. Returns are normally distributed with annual volatility 15% and expected return of 10%. Fund failures arrive randomly following a log-logistic distribution. Upon failure, NAV drops 50% and the investor receives the remaining assets as a liquidating dividend. Panel A shows results for new funds. Panel B shows results for funds with initial age of 24 months.

Panel A. Age = 0						
Notice	Lockup					
	12	24	36	48	60	
1	0.03	0.12	0.59	1.39	2.27	
2	0.06	0.16	0.67	1.47	2.34	
3	0.09	0.21	0.74	1.56	2.42	
4	0.12	0.26	0.81	1.64	2.50	
5	0.15	0.31	0.89	1.73	2.57	

Panel B. Age = 24						
Notice	Lockup					
	12	24	36	48	60	
1	0.14	0.58	1.18	1.75	2.24	
2	0.19	0.65	1.25	1.81	2.28	
3	0.25	0.71	1.31	1.86	2.32	
4	0.30	0.78	1.37	1.92	2.36	
5	0.36	0.85	1.44	1.97	2.39	

Table 6. Hedge Fund Values as a Function of Fund Age and Horizon.

Listed in Panel A is the value of a hedge fund with an unrestricted liquidity option over a range of initial ages and horizons, T , both in months. Panel B shows values when the liquidity option is restricted by a lockup of two years and notice period of three months. Panel C shows values when no liquidity option exists. The hedge fund has initial NAV of \$100. Returns are normally distributed with annual volatility 15% and expected return of 10%. Fund failures arrive randomly following a log-logistic distribution. Upon failure, NAV drops 50% and the investor receives the remaining assets as a liquidating dividend.

Panel A. Unrestricted Liquidity Option, $O_{0,1}$					
Horizon	Age				
	0	24	48	72	96
72	105.69	100.65	100.04	100.68	102.86
96	105.84	100.80	100.27	101.77	104.81
120	106.04	101.06	100.80	103.32	107.09
144	106.34	101.47	101.71	105.21	109.47
168	106.76	102.09	102.99	107.23	111.98

Panel B. Restricted Liquidity Option, $O_{0,1}^{L,NP}$					
Horizon	Age				
	0	24	48	72	96
72	105.50	99.84	99.00	100.45	102.85
96	105.64	100.03	99.49	101.67	104.81
120	105.83	100.35	100.27	103.28	107.08
144	106.12	100.84	101.39	105.19	109.47
168	106.52	101.56	102.81	107.23	111.98

Panel C. No Liquidity Option, $H_{0,1}$					
Horizon	Age				
	0	24	48	72	96
72	102.56	97.56	97.89	100.23	102.83
96	101.53	97.44	98.48	101.53	104.80
120	101.32	97.90	99.53	103.21	107.08
144	101.69	98.78	100.91	105.16	109.47
168	102.47	99.95	102.54	107.22	111.98

Table 7. Hedge Fund Values.

Listed in Panel A are values per share of a new hedge fund with initial NAV of \$100 and a ten-year life. Returns are normally distributed with annual volatility of 15% and expected return as listed. Fund failures arrive randomly following a log-logistic distribution. Upon failure, NAV drops 50% and the investor receives the remaining assets as a liquidating dividend. The five columns are: the value to an investor with no liquidity option, $H_{0,1}$; the value to an investor with a liquidity option subject to a two-year lockup and a three-month notice period, $O_{0,1}^{L,NP}$; the value subject only to a lockup, $O_{0,1}^L$; the value subject only to a notice period, $O_{0,1}^{NP}$; and the value when no restrictions are in place, $O_{0,1}$. Panel B lists values when the fund is 24 months old at the time of investment and has cumulative return at that time equal to the cross-sectional mean.

Panel A. Age = 0					
μ	$H_{0,1}$	$O_{0,1}^{L,NP}$	$O_{0,1}^L$	$O_{0,1}^{NP}$	$O_{0,1}$
6%	72.59	95.78	96.69	100.27	100.29
7%	78.63	97.99	98.68	100.79	100.81
8%	85.37	100.26	100.71	101.69	101.71
9%	92.90	102.63	102.87	103.16	103.21
10%	101.32	105.83	105.96	105.95	106.04
11%	110.72	111.89	111.98	111.91	111.99
12%	121.23	121.39	121.41	121.39	121.41
13%	132.98	133.00	133.00	133.00	133.00
14%	146.13	146.13	146.13	146.13	146.13

Panel B. Age = 24					
μ	$H_{0,1}$	$O_{0,1}^{L,NP}$	$O_{0,1}^L$	$O_{0,1}^{NP}$	$O_{0,1}$
6%	71.72	91.27	92.29	99.16	100.00
7%	77.26	93.29	94.11	99.41	100.00
8%	83.42	95.36	95.97	99.65	100.00
9%	90.28	97.48	97.88	99.90	100.00
10%	97.90	100.35	100.55	100.93	101.06
11%	106.39	106.67	106.73	106.70	106.75
12%	115.85	115.86	115.86	115.86	115.86
13%	126.39	126.39	126.39	126.39	126.39
14%	138.13	138.13	138.13	138.13	138.13

Figure 1. The Lattice.

The figure depicts the binomial lattice structure representing the normal regime of a hedge fund. The passage of time is represented by moving from left to right and denoted by an increase in the variable t . Horizontal movement represents a positive return. Diagonal downward movement represents a negative return. We refer to the combination of time step t and level j as node (t,j) .





Figure 2. Durations of Defunct Funds and History Lengths of Live Funds.

Panel A shows the percentage of the 4,260 defunct funds in the 2005 CISDM database with duration equal to the value on the horizontal axis. Panel B shows the percentage of the 4,272 live funds in the 2005 CISDM database with history lengths equal to the value on the horizontal axis.

Panel A. Durations of Defunct Funds



Panel B. History Lengths of Live Funds




Figure 3. Log-logistic Function.

Depicted by hollow squares is the number of defunct hedge funds in the 2005 CISDM hedge fund database with lifespan equal to the values on the horizontal axis. Depicted in bold is the predicted number of hedge funds with lifespan equal to the values on the horizontal axis. Predicted number of hedge funds at lifespan t equals the hazard rate of the log-logistic function fitted to the data evaluated at t times the number of funds in the database with lifespan greater than or equal to t . Data include 8,532 funds with data through 2005.




Figure 4. Probability of Failure in the Lattice.

Figure shows the probability of hedge fund failure as a function of age:

$$\lambda q (\lambda t)^{q-1} / [1 + (\lambda t)^q]$$

where $\lambda = 0.0129$ and $q = 1.6517$ are parameters of the log-logistic function estimated from the durations of hedge funds in the 2005 CISDM database and t is the age of the fund, in monthly time steps.





Figure 5. Probability of Failure in the Lattice with Performance Covariate.

Panels show the probability of hedge fund failure as a function of age and return for a fund with annual expected return of 10% and volatility as indicated. In the figures, age ranges from 1 to 36 months and return is measured by the number of down steps taken in a binomial lattice. Probability of failure at a given node is:

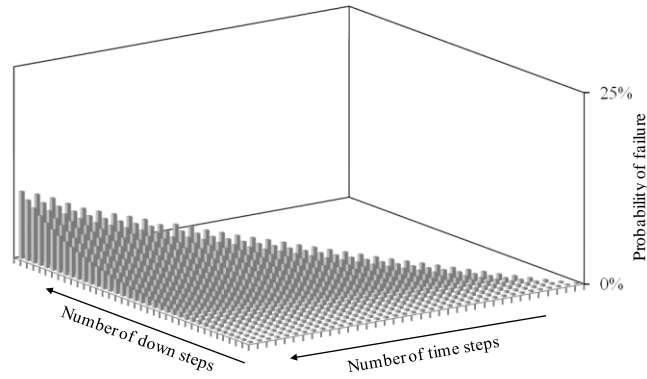
$$\lambda q (\lambda t)^{q-1} / [1 + (\lambda t)^q] \exp(z\beta)$$

where $\lambda = 0.0129$ and $q = 1.6517$ are parameters of the log-logistic function estimated from the durations of hedge funds in the 2005 CISDM database, t is the age of the fund, $\beta = -0.3237$ is estimated from the 2005 CISDM database, and z is the return of the fund at a node expressed as the number of cross-sectional standard deviations from the cross-sectional mean, both estimated from the 2005 CISDM database.

Panel A. Volatility = 15%



Panel B. Volatility = 25%



Panel C. Volatility = 40%





Figure 6. Liquidity Option Values and Restriction Costs.

Listed in Panel A are unrestricted liquidity option values (left graph) and the combined cost of a two-year lockup and a three-month notice period, for a fund with initial NAV of \$100 and a ten-year life. Returns are normally distributed with annual volatility of 15% and expected return as listed. Fund failures arrive randomly following a log-logistic distribution. The λ parameter estimated from the data is scaled by the factor listed. Upon failure, the investor receives 50% of the remaining assets as a liquidating dividend. Panel B shows option values and restriction costs when the percentage of assets recovered upon failure, l , are varied as listed.

Panel A. Impact of Hazard Rate



Panel B. Impact of Loss

