

Limits of Limits of Arbitrage: Theory and Evidence*

Johan Hombert[†] David Thesmar[‡]

March 2, 2009

Abstract

We present a model where arbitrageurs operate on an asset market that can be hit by information shocks. Before entering the market, arbitrageurs are allowed to optimize their capital structure, in order to take advantage of potential underpricing. We find that, at equilibrium, some arbitrageurs always receive funding, even in low information environments. Other arbitrageurs only receive funding in high information environments. The model makes two easily testable predictions: first, arbitrageurs with stable funding should experience more mean-reversion in returns, in particular following low performance. Second, this larger mean-reversion should be lower, if many other funds have stable fundings. We test these predictions on a sample of hedge funds, some of which impose impediments to withdrawal to their investors.

*This paper has benefited from financial support from the BNP Paribas Hedge Fund Centre at HEC, and the Paul Wooley Centre for the Study of Market Dysfunctionalities. We are grateful to Serge Darolles, Vincent Poudroux, Isabelle Serot, Guillaume Simon at SGAM - AI for decisive help with hedge fund data.

[†]ENSAE and GREMAQ. hombert@ensae.fr

[‡]HEC, CEPR and ECGI. thesmar@hec.fr

1 Introduction

In the literature on limits to arbitrage, a widening of the mispricing of an asset may lead arbitrageurs to unwind their positions, which further amplifies the initial mispricing (Shleifer and Vishny, 1997, Gromb and Vayanos, 2002). Such forced unwinding occurs because, as arbitrageurs lose money on their trades, their investors (brokers, banks, limited partners etc) demand early reimbursement of their claims. Thus, existing theories of the limits to arbitrage assume that arbitrageurs cannot design their capital structure ex ante (for instance, by taking on long term debt) in order to avoid such value destroying events.

This paper starts from the simple fact that this assumption does not always hold in reality, and investigates its theoretical and empirical consequences. In the hedge fund industry, investors often agree to limit their ability to withdraw their funds. About 20% of the hedge funds in our sample have lock up periods of typically one, or even two years, during which investors cannot redeem their shares (Aragon, 2007, has a similar proportion). Once they are able to do so, they must give the fund advance notice (typically a month) and then obtain redemption at fixed dates (typically a quarter). For the average hedge fund in our sample, we estimate the minimum duration of funds to be equal to 5 months, and 10 months for funds with lock-up periods. Interestingly, such share restrictions can be found with hedge funds investing in illiquid securities (such as fixed income), but also with funds dealing with stocks (such as “long short equity” funds). They are what we call “limits of limits of arbitrage”: thanks to them, some market participants can afford to underperform in the short run while they hold on to ultimately profitable arbitrage opportunities.

Thus, at least some arbitrageurs choose the maturity of their investors’ claims. To under-

stand the determinants and consequences of such a capital structure decision, we first build a model where arbitrageurs optimally design the securities that they issue, and then engage in arbitrage on the same market. Arbitrageurs differ in skill. We posit that arbitrageur skill affects long term asset payoffs in some states of nature only (“low information states”). We first find that, at equilibrium, prices in low information states are lower, because scarce arbitraging skills are needed to trade in these states. Furthermore, at equilibrium investors guarantee funding to skilled arbitrageurs in low information (low price) states, while unskilled arbitrageurs only receive funding in high information (high price) states. The intuition comes from the asset price equilibrium: if investors did not guarantee funds to some arbitrageurs in low information states, asset prices in these states would collapse, which would make investment attractive. Even when arbitrageur skill is not contractible upon, the equilibrium capital structure choice is separating: skilled arbitrageurs choose guaranteed funding, while unskilled arbitrageurs choose funding only contingent on high (high prices) information. Thus, there is optimal differentiation at equilibrium.

Our model generates two easily testable predictions. First, conditional on past bad performance, funds with guaranteed funding outperform other funds. As argued above, in low information states, scarce skills are needed which lowers current prices. Thus, funds with guaranteed funding invest more often in states where the assets are underpriced, and thus outperform other funds who take less advantage of underpricing. Our second prediction is cross sectional: in industries where the fraction of arbitrageurs with guaranteed funding is large, these arbitrageurs overperform other funds less. The mechanism is deeply rooted in the model: if more arbitrageurs receive guaranteed funding, underpricing in the low information state will be reduced. Since these arbitrageurs benefit from underpricing, their

overperformance will be reduced.

We then test these two predictions on hedge fund data. We use the fact that, in our data, funds with impediments to withdrawal (such as long redemption periods, or lock-up periods) experience less outflows when they underperform (Ding et al, 2008, find similar evidence in a smaller sample). Thus, these share restrictions, which we can observe from the data, are a good proxy for “guaranteed funding” in our model. We find that the first prediction of our model holds in the data: conditional on bad past performance, funds with impediments to withdrawal do “bounce back more”, i.e. have higher expected returns. We also find some support for the second prediction. To test it, we look at investment styles where impediments to withdrawal are prevalent. We find that, in these styles, funds with such share restrictions overperform other funds relatively less than in styles where such impediments are relatively rare, although these results are somewhat less robust.

This paper contributes to two strands of literature. First, we extend Shleifer and Vishny’s model of limits of arbitrage by allowing arbitrageurs to optimally choose their capital structure in order to avoid inefficient liquidation. In this sense, our paper is closely related to independent work by Stein (2009), which is the only paper, to the best of our knowledge, that explicitly seeks to endogenize arbitrageurs’ capital structures. Compared to his model, our theory endogenizes the cost of external finance more explicitly and makes testable predictions on observed arbitrageur returns, that we can bring to the data. What is common to the two models is that underpricing in bad states of nature leads to more investment in these states, through the optimal structure choice. This feedback mechanism is not present in Shleifer and Vishny (1997) nor in Stein (2005), who does not endogeneize prices. The rest of the limits to arbitrage literature considers the destabilizing feedback that goes through

the wealth of arbitrageurs. Gromb and Vayanos (2002), Acharya and Viswanathan (2007) and Brunnermeier and Pedersen (2008) model intermediaries that need to unwind their positions when collateral prices decrease, which amplifies price drops. In all these models, even if mispricing can be very large, there is no contractual way to take advantage of this.¹

Second, we shed new light on existing evidence from the (mostly hedge) funds literature. First, our findings are related to a recent paper on open end mutual funds by Coval and Stafford (2008): they look at asset fire sales following massive redemptions at mutual funds, and find a significant price impact. Thus, their paper suggests (but does not test) that mutual fund performance should display some persistence, in particular conditional on past low performance. We propose a theory why some funds may seek protection against massive redemptions, and what returns dynamics should look like in protected and unprotected funds. Second, some papers show that the presence of impediments to withdrawal is correlated with unconditional fund performance (Aragon, 2007, Agarwal, Daniel and Naik, 2008): their explanation is that investors earn a premium for the illiquidity of their investment. Other

¹Also related to this paper is Lerner and Schoar (2004). They test a model where (private equity) fund managers make their shares illiquid in order to select “patient” investors. Their capital structure focus differs from ours in an important way: they look at the ability to sell shares to *other investors*, while we look at the ability to sell shares to the *fund*. In addition, if we were to transpose their mechanism in our setting where funds interact through buying and selling the same asset, funds would screen “patient investors” until mispricing disappears. There would be no prediction on the link between share restriction and fund returns dynamics.

Casamatta and Pouget (2008) solve a model where investors give fund managers incentives to search for information on assets. In their model, the optimal contract features short term performance pay. The cost of such contracts is that they reduce market efficiency by deterring information acquisition.

papers have informally argued that hedge funds act as liquidity providers (see e.g. Agarwal, Fung, Loon and Naik, 2008). The present article suggests a potential reason why illiquid funds can afford to issue illiquid shares in the first place: because illiquidity allows them to reap the gains of arbitrage, they can pay the illiquidity premium to their investors. Third, we develop and test a theory of the mean reversion of fund returns. Interestingly, some of the existing hedge funds literature has focused on the *positive* relation between autocorrelation and share restrictions (see e.g. Aragon, 2007) while we find clear evidence of such a *negative* relation. The difference between these studies and ours is the frequency at which autocorrelation is computed: we work at the annual level, while existing papers work at the monthly level. At the monthly frequency, the existing literature argues that reported returns of illiquid assets are smoothed. At the annual frequency, this paper argues that arbitrage induces a mean reversion in fund returns. To some extent, such evidence is reminiscent from insights from the strategic allocation literature (Campbell and Viceira, 2002) which argues that long term investors have a comparative advantage at investing in mean reverting assets.

The rest of the paper follows a simple structure. Section 2 describes, solves the model and derives predictions and comparative statics. Section 3 tests the model. Section 4 concludes.

2 Model

2.1 Set-Up

This framework borrows from Hombert (2007)'s model of fire sales in equilibrium. There are competitive, risk neutral, investors. Investors want to purchase an asset which is in

unit supply, but cannot do so themselves. They delegate this task to a measure 1 of fund managers. Fund managers are risk neutral and limitedly liable; each of them starts with initial wealth A .

2.1.1 Sequence of Events

There are four periods $t = 0, 1, 2, 3$ and the discount rate is zero. At $t = 0$, investors contract with managers. The optimal contract will specify the amount of funds that the investor will entrust to the manager, both in $t = 1$ and 2, and conditional on the state of nature in $t = 2$ (see below).

At date $t = 1$, each fund manager *learns about the asset* he will be trading: the acquired knowledge will only be useful in period $t = 2$ (see below). Learning effort costs C to the manager. We assume here that learning effort is not contractible but this is not necessary (more on this below). With high learning effort, the manager becomes skilled with probability μ ; with low learning effort, with probability $\mu - \Delta\mu$. The manager does not know whether he is skilled until period 3. After the learning phase, managers use entrusted funds to purchase assets at unit price P . The market for assets clears.

At date $t = 2$, the market can be in one of three states. With probability $1 - \lambda - \varepsilon$, the market is in state U : in this state, knowledge acquired in period 1 is useless (think for instance of a bull market where everyone can generate high returns). It becomes public knowledge that the asset will generate $t = 3$ cash flows of $V > 0$. All fund managers liquidate their positions from $t = 1$, pay off their investors, and use newly entrusted funds (as specified in the optimal contract) to repurchase the same assets. The market clears again at price P_U .

With probability λ , the market is in state M . In this state, we assume that a second

asset, which is a priori not distinguishable from the first asset, appears. Because they cannot be differentiated from each other, both assets trade at the same price, but we assume that the second asset has zero present value, while the first asset has, as in state U , a PV of V . Furthermore, we assume that only a fraction μ of the managers picks the “right” asset. The important hypothesis is that, in state M , the asset PV does not depend on $t = 1$ effort. Thus, compared to state U , state M is a *bad* state, in the sense that there is less information than in state U , but the state is equally bad for all managers, irrespective of their $t = 1$ learning decisions. In this state, the market for assets clears at price P_M .

Last, with probability ε , the market is in state D . Exactly as in state M , a second asset appears that has a PV of zero, but this time skilled managers can differentiate between the two. Thus, an important difference between states M and D is that in state D , date 1 learning effort matters. In this sense, state D also is a *bad* state, but it is worse for managers who did not learn in $t = 1$. In this state, the “right” asset market clears at price P_D (which is also the price of the “wrong” asset, whose market we do not model).

At date $t = 3$, assets held in portfolios mature and payoffs are realized. If the “right” asset is held, its payoff is V . In states M and D , we assume that only $V - B$ can be pledged to the investors. We think of B as the rent of an unmodelled agency conflict in period 3: for instance, the manager can sell the asset on a black market for price B and consume the proceeds. To simplify exposition, we assume that this agency conflict does not exist in state U (in which case the entire present value of the asset V can be pledged to investors). All intuitions of this model would carry through without this assumption.

All in all, states U , M and D vary along two key dimensions. First, in state U , expected cash flows from assets are higher than in states M and D . Expected present value in U is

always V ; in state M it is only μV . In state D only skilled managers will be able to buy good assets, so the expected payoff is at most μV . This feature of the model ($\mu < 1$) is not entirely necessary for most intuitions to carry through; we will explain why later on. The second difference between the three states is that, in state D , managerial skill matters more. Thus, a manager who is committed more funds in state D will have more incentive to learn. This second dimension of our model is essential.

2.1.2 Contracts

We assume that the financial contract specifies four amounts entrusted to the manager: I , in period 1, and (I_U, I_M, I_D) in period 2, conditional on states U , M and D . Thus, we make two implicit (and mostly simplifying) assumptions. First, we assume that date 1 learning effort cannot be contracted upon. This assumption is not necessary to obtain our results, but it simplifies calculations (more on this later). Second, we assume that the date 2 state of nature is contractible. This could be the case for instance if period 2 returns were contractible. It is precisely the goal of this paper to study the impact of the contingent financing of arbitrageurs.²

Given these assumptions on verifiability, there is no need to increase the contract space in our model: the four instruments I , I_U , I_M and I_D are sufficient to reach the second best optimum. For instance, assume the investors commits to paying a given amount to the manager in case of success (positive asset pay off) at date $t = 3$ in order to induce learning effort at date $t = 1$. As it turns out, such transfer would only consume part of the income

²An alternative to making $t = 2$ state of nature contractible could be to implement contingent control rights allocation à la Aghion and Bolton (1992).

pledgeable to the investor, so I_U , I_M or I_D would have to decrease. This would reduce the assets under management, and hence the NPV of the fund. Thus, adding an incentive fee on top of the private benefit B per unit of asset would not be part of an optimal contract.

In the empirical application, we will think of contracts where $I_D > 0$ as contracts with impediments to withdrawals. Such contracts guarantee $t = 2$ inflows even when the fund underperforms (asset prices go down in state D).

2.1.3 Modelling Strategy

Before we solve the model, it is worthwhile to discuss our modelling strategy. With moral hazard, another possible strategy could have been to think of withdrawals as an ex ante optimal “punishment” strategy. Underperforming managers are punished for low effort provision, while well performing managers are rewarded through continuation. Such a model delivers similar comparative static properties as the one we study in this paper, but the implied contract has the important drawback of not being renegotiation-proof. Once low effort has been provided, assets are fairly priced in equilibrium, and their expected return is non negative. Thus, shutting down the fund is never an optimal decision ex post and the punishment is non credible.

One second alternative would have been to model limits of arbitrage as arising because investors learn about the fund manager’s skill, assuming such a skill is fixed from the beginning (i.e. not obtained through learning). If the fund underperforms, then it becomes optimal to withdraw investment because the chances that the manager is incompetent are high. In such a model, there would be no reason for an investor to lock his money up in the fund because there is no efficiency gain to do so.

To make impediments to withdrawals optimal from a contracting perspective, they must have an incentive property, so we choose a moral hazard setting (learning entails an “effort”). An alternative model would be to assume that managers have fixed types (skilled or unskilled) and that investors seek to design separating contracts. Such a model would be almost identical to the model we present in this paper, except that learning is exogenous. Such a model would generate identical predictions to the ones we derive and test here.

2.2 Solving the Model

We first derive the optimal contracts for given expected asset prices, and then solve for the rational expectations equilibrium of the asset market. This allows us to (1) characterize the equilibrium and (2) find a relationship between impediments to withdrawals ($I_D > 0$), and the equilibrium returns of the funds.

2.2.1 Optimal Contracts

In this Section, we take the sequence of future prices P , P_U , P_M and P_D as given. The optimal contract solves the manager’s objective function, which is the project’s NPV, under the constraints that the profit pledgeable to investors is nonnegative and that the manager exerts the desired level of learning effort (Tirole, 2006).

We first focus on high learning effort funds:

$$\begin{aligned}
& \max_{I, I_U, I_M, I_D} \left\{ \begin{aligned} & I \cdot [(1 - \lambda - \varepsilon) P_U + \lambda P_M + \varepsilon P_D - P] \\ & + (1 - \lambda - \varepsilon) I_U \cdot [V - P_U] \\ & + \lambda I_M \cdot [\mu V - P_M] + \varepsilon I_D \cdot [\mu V - P_D] \end{aligned} \right\} \\
& \text{s.t.} \quad I \cdot [(1 - \lambda - \varepsilon) P_U + \lambda P_M + \varepsilon P_D - P] + (1 - \lambda - \varepsilon) I_U \cdot [V - P_U] \\
& \quad + \lambda I_M \cdot [\mu(V - B) - P_M] + \varepsilon I_D \cdot [\mu(V - B) - P_D] + A \geq 0 \\
& \quad \varepsilon \Delta \mu \cdot B \cdot I_D > C
\end{aligned}$$

The objective function is the overall NPV of the fund. The first term is the total profit made between period 1 and 2, which is equal to the expected price increase times the amount invested in $t = 1$. Given that this profit is free from any agency consideration, it can be pledged to the investor at 100%, which is why it also appears as such in the first (investor participation) constraint. The second term is the $t = 2$ NPV realized in state U , which can also be fully pledged. The third term is the expected NPV in state M . In this case, the manager will purchase the right asset with probability μ (since learning effort has been made) and, as appears in the first constraint, only $\mu(V - B)$ per asset purchased can be pledged at $t = 0$. In state D , the conditional expected payoff per asset is the same, because the manager puts in high effort. The second constraint is the manager's incentive compatibility constraint which ensures that, in period 1, high learning effort is always preferred. Given our parameters restrictions below, this constraint will never binds at equilibrium.

It is clear from the above problem that:

$$P = (1 - \lambda - \varepsilon) P_U + \lambda P_M + \varepsilon P_D$$

will have to hold in equilibrium. If this is not the case, I will be equal to $+\infty$ or $-\infty$.

Hence, markets in $t = 1$ are fully efficient in this model because there is no agency friction in $t = 1$: any profit from arbitrage is pledgeable, so that infinite amount of funds can be used to finance arbitrageurs. This reduces arbitrage opportunities to zero. For the same reason, the same happens in state U : $P_U = V$. Thus, all funds receive an indeterminate amount of funding in state U .

Moreover, it is easy to see that $P_M \leq \mu V$ and $P_D \leq \mu V$ have to hold in equilibrium, otherwise no fund would be willing to hold the asset in state M or in state D . At the same time, $P_M > \mu(V - B)$ and $P_D > \mu(V - B)$. This comes from the fact that the marginal pledgeable income of investment has to be strictly negative in equilibrium. If this is not the case, fund managers could raise money to invest more, as the NPV of doing so is strictly positive. This would contradict the equilibrium.

Given these properties and the convenient linearity of the problem, we obtain that:

$$\begin{aligned} I_M &= \frac{1}{\lambda} \frac{A}{P_M - \mu(V - B)}, \\ I_D &= 0, \\ NPV &= \frac{A}{P_M - \mu(V - B)} \cdot \mu(V - P_M) - C, \end{aligned}$$

if $P_M \leq P_D$. Hence, if the price in state M is low enough compared to the price in state D , it is then efficient (in terms of NPV) to allocate all pledgeable income in state M where the asset is relatively cheap. In contrast, as soon as $P_M \geq P_D$:

$$\begin{aligned} I_U &= 0, \\ I_D &= \frac{1}{\varepsilon} \frac{A}{P_D - \mu(V - B)}, \\ NPV &= \frac{A}{P_D - \mu(V - B)} (\mu V - P_D) - C. \end{aligned}$$

When the asset is cheap in state D , it is efficient to allocate all pledgeable income in state D .

Computations and expressions are very similar when the optimal learning effort is low. In this case, the fund will invest only in state D if and only if $P_M > \frac{\mu}{\mu - \Delta\mu} P_D$. This leads us to the following lemma:

Lemma 1 *For given asset prices $P_M \leq V$ and $P_D \leq \mu V$, there are five regimes:*

In all regimes, all funds receive funding in state U . In addition.

1. $P_M < P_D$. *In this case, both high and low learning effort funds invest only in state M .*
2. $P_M = P_D$. *High effort funds are indifferent between investing in state M and D . Low effort funds invest only in state M .*
3. $P_D < P_M < \frac{\mu}{\mu - \Delta\mu} P_D$. *Then, high effort funds only invest in state D , low effort funds only in state M .*
4. $P_M = \frac{\mu}{\mu - \Delta\mu} P_D$. *High effort funds invest in state D and low effort funds are indifferent between investing in state M and D .*
5. $\frac{\mu}{\mu - \Delta\mu} P_D < P_M$. *Both high and low effort funds only invest in state D .*

The results of this lemma are intuitive: high P_M discourages funds to invest in state M . In addition, high effort funds have higher returns to investing in state D , since this is when learning effort pays off. Hence, high effort funds are ready to invest in state D for higher levels of P_D (i.e. lower expected returns).

The above lemma also indicates that cases 1 and 5 cannot be equilibrium outcomes, since in these cases there is no demand for assets in either state D or M . Putting aside the knife-edge cases 2 and 4, this suggests that in equilibrium both levels of learning effort coexist: high effort funds only invest in state D , while low effort funds invest in M only.³ We now turn to the description of the equilibrium.

2.2.2 Equilibrium

Following the discussion above, we restrict ourselves to $P_D \leq P_M \leq \frac{\mu}{\mu - \Delta\mu} P_D$. Let α be the equilibrium fraction of high effort funds. In equilibrium, since both categories of funds coexist, funds have to be indifferent, ex ante, between putting in high effort (and buy in state D) or low effort (and buy in state M):

$$\frac{\mu V - P_D}{P_D - \mu(V - B)} - \frac{C}{A} = \frac{\mu V - P_M}{P_M - \mu(V - B)}. \quad (1)$$

We now need to compute equilibrium prices P_M and P_D . Aggregate asset demand by funds in state M and state D has to be equal to supply (assumed equal to 1). Hence:

$$P_M = \mu(V - B) + \frac{\mu(1 - \alpha)A}{\lambda} \quad (2)$$

$$P_D = \mu(V - B) + \frac{\mu\alpha A}{\varepsilon}. \quad (3)$$

The price in each state is higher, the higher the expected payoff, the higher the equity of managers, and the higher the number of funds operating in this state. Plugging back (2)

³Case 2 cannot actually arise in equilibrium, otherwise a high training effort fund would be indifferent in $t = 2$ between investing in state M or D . In $t = 1$, it would then be optimal to make low effort and invest in state M only, to save the training effort cost, hence there would be not demand for the asset in state M . By contrast, case 4 can be an equilibrium outcome for some parameter values. To clarify exposition, we rule them out in the following.

and (3) into indifference condition (1), we obtain the following equation for the equilibrium α :

$$\frac{\varepsilon}{\alpha} - \frac{C}{B} = \frac{\lambda}{1 - \alpha}. \quad (4)$$

It is straightforward to see that $\alpha \in]0; 1[$. Moreover, α is increasing in ε , and decreasing in λ and C/B . When the cost of making effort (C) decreases, or when the gains of making effort (ε) are larger, there will be more high effort funds operating in equilibrium.

So far we have assumed that, once the contract is signed, the fund manager puts in the expected amount of effort. It is straightforward to see that a manager with $I_D = 0$ will make no learning effort, since it will never pay off. A fund manager with $I_D > 0$ puts in high effort if and only if his incentive constraint is satisfied:

$$I_D = \frac{1}{\alpha\mu} > \frac{C}{\varepsilon B \Delta\mu}. \quad (5)$$

i.e. I_D is large enough to make the gain of learning $\varepsilon \Delta\mu B \cdot I_D$ larger than the effort cost C .

Finally, we need to ensure that asset prices in period 2 are below their fundamental value in equilibrium (else conditions (2) and (3) do not apply). This occurs if and only if:

$$A < \frac{\lambda B}{1 - \alpha}. \quad (6)$$

Intuitively, if fund managers have little equity, their demand will be so low that prices not reach their fundamental values, even in state M .

Hence, an equilibrium is defined by equation (4), under conditions (5) and (6). Equilibrium prices also have to satisfy $P_D \leq P_M \leq \frac{\mu}{\mu - \Delta\mu} P_D$. The following proposition characterises such an equilibrium, and provides a parameter condition for its existence:

Proposition 2 *Equilibrium Characterization*

There exists \bar{A} and $\underline{\Delta\mu} < \mu$ such that, if:

$$\Delta\mu > \underline{\Delta\mu} \text{ and } A < \bar{A}$$

1. The only equilibrium is an equilibrium where $\alpha \in]0; 1[$ funds make learning effort and are only committed funds in states U and D , and $1 - \alpha$ funds make low learning effort and are only entrusted funds in states U and M .
2. α is defined by equation (4). Equilibrium prices are such that $P_M > P_D$.
3. The ex ante optimal contract is renegotiation-proof in equilibrium.

Proof. Let α^* be the (unique) positive solution of equation (4). Let $\bar{A} = \lambda B / (1 - \alpha^*)$. The condition $A < \bar{A}$ is equivalent to condition (6) which is therefore satisfied. Let $\underline{\Delta\mu} = C\alpha^*\mu/\varepsilon B$: this ensures that the incentive compatibility constraint for high effort funds holds. From (4), it is easy to see that $\underline{\Delta\mu} < \mu$.

From equilibrium prices (2) and (3), it is easy to obtain that:

$$P_D - P_M = \mu A \cdot \left(\frac{\alpha^*}{\varepsilon} - \frac{1 - \alpha^*}{\lambda} \right) < 0$$

which is negative by virtue of (4). QED ■

The optimal contract is renegotiation-proof because, in equilibrium, the asset price is always above the marginal pledgeable payoff, but below the marginal NPV. As a result, the manager cannot raise new funds (he can only promise a negative income $V - B - P_i$, $i = M, D$) nor is he willing to cut down investment (he obtains utility B per asset invested). Put differently, the contract is renegotiation proof because continuation is as optimal ex post as it is ex ante: the size of the surplus does not increase nor decrease and there is thus no scope for renegotiation.

One can compute the relative underpricing in state D as the difference $P_M - P_D$. Given that the expected PV of the assets is $\mu(V - B)$ in both states, $P_M - P_D$ measures the price difference that is not due to a difference in expected payoff, but simply a lack of invested funds. This underpricing is given by:

$$P_M - P_D = \left(\frac{1 - \alpha}{\lambda} - \frac{\alpha}{\varepsilon} \right) \mu A, \quad (7)$$

therefore it is an increasing function of C/B . It is equal to 0 if $C/B = 0$. As the cost of learning tends to zero, more and more funds are willing to raise money in state D . This brings prices in state D closer to fundamental value.

This remark extends the model of Shleifer and Vishny (1997) to a full contracting framework, where investors are also able to commit funding in the low state of nature (i.e. when the asset is underpriced). What we show is that, when assets are underpriced, there is an incentive for another class of fund to specialize in this state. Yet, because arbitrage in this state is costly (managers need to learn enough about the asset), state D prices cannot increase too much.

A related point is that learning is necessary in our model to obtain underpricing. Assume that $\Delta\mu < \underline{\Delta\mu}$, so that it is never optimal for funds with positive inflows in state D to learn in period 1. In this case, we are looking for an equilibrium where both funds investing in states M , and in state D , make no learning effort. In this case, it is easy to verify that there is no price distortion. The intuition is that entering state D is now costless as it entails no learning effort: free entry in the two states makes returns identical.

2.3 Predictions of the Model

Our first prediction is related to net of fee conditional returns. For both types of funds, expected returns conditional on good performance in period 2 (i.e. in state U) are equal to zero, since $P_U = V$. This comes from the fact that there is no agency friction B in this state. Including one would not change the results: both types of funds would still have the same expected returns in this state, since they would purchase the same asset at the same price.

Our model has more interesting predictions on returns conditional on low performance in period 2:

$$E(R_3|R_2 \text{ is low}, I_D > 0) = \mu(V - B) - P_D = -\frac{\mu\alpha A}{\varepsilon}, \quad (8)$$

$$E(R_3|R_2 \text{ is low}, I_D = 0) = \mu(V - B) - P_M = -\frac{\mu(1 - \alpha)A}{\lambda}. \quad (9)$$

It appears clearly that expected returns of high effort funds are larger than expected returns of low effort funds, since high effort funds invest in state D , where assets are significantly underpriced ($P_D < P_M$). What is interesting is that this prediction holds in equilibrium, even though “entry” in both states of nature is free at the contracting stage.

Prediction 1 *High effort funds exhibit more mean reversion in returns, in particular when past returns are low. More precisely:*

1. *Conditional on high past returns, both funds have similar expected returns.*

$$E(R_3|R_2 \text{ is high}, I_D > 0) = E(R_3|R_2 \text{ is high}, I_D = 0)$$

2. *Conditional on low past returns, high effort funds overperform low effort ones.*

$$E(R_3|R_2 \text{ is low}, I_D > 0) > E(R_3|R_2 \text{ is low}, I_D = 0)$$

Before we proceed, we note that the model would have exactly the same properties (under slightly weaker conditions on parameters) if learning effort was contractible. Thus, the differential mean reversion does not hinge on the contractibility or incontractibility of learning effort; it is just slightly easier to obtain under full contracting.

The above prediction is related to Aragon (2007) and Agarwal et al (2008), who find empirically that hedge funds with impediments to withdrawal tend to exhibit superior performance (even after controlling for usual risk factors). They interpret this correlation as evidence that investors demand a premium for holding illiquid (i.e. locked up) shares.⁴ And indeed, given the loss of (put) option value, the cost of illiquidity to investors can be quite sizeable (Ang and Bollen, 2008, perform a calibration using a real option model). Our model has the feature that high effort funds may exhibit higher performance under some circumstances. Using (8) and (9), we find easily that the excess unconditional performance of high effort funds is given by:

$$E(R_3|\text{high effort}) - E(R_3|\text{low effort}) = \mu(1 - 2\alpha) A$$

High effort funds outperform in our model as long as $\alpha < 1/2$. If α is small enough, fewer funds invest in state D while more funds invest in state M . Thus, underpricing in state D is large enough to make high effort funds outperform low effort ones.

Our model does not only predict that high effort fund' returns mean revert more, but

⁴It is interesting to note that the presence of impediments to withdrawals does not necessarily mean that investment is illiquid. One possibility is that shares, even though not immediately redeemable, can be traded among investors on a secondary market (for the description of such a market, see for instance Ramadorai, 2008).

also makes predictions on the relative size of the mean reversion. Intuitively, as the share of high effort funds α increases, there is more investment in state D , which increases P_D and reduces P_M . Given that high effort funds invest in state D , their outperformance should therefore be reduced. Thus, across asset markets, α and the overperformance of high effort funds conditional on past low returns should be negatively correlated. Given that α and the extent of mean-reversion are both endogenous, we need to verify that this intuition holds in the model. We do this in the proposition below:

Prediction 2 *When ε increases, when λ decreases, when C decreases, when B increases,*

1. α increases
2. *The outperformance of high effort funds, conditional on bad performance:*

$$\rho = E(R_3|R_2 \text{ is low}, I_D > 0) - E(R_3|R_2 \text{ is low}, I_D = 0)$$

decreases.

Proof. From (8) and (9):

$$\rho = -\mu A \cdot \left(\frac{\alpha}{\varepsilon} - \frac{1 - \alpha}{\lambda} \right)$$

which is a decreasing function of α .

First, as C/B decreases, α increases, which reduces ρ .

Second, as ε increase, α increases (see equation 4). Using (4), the outperformance ρ can be rewritten as:

$$\rho = -\mu A \cdot \left(\frac{1}{\lambda/(1 - \alpha) + C/B} - \frac{1 - \alpha}{\lambda} \right)$$

which is decreasing in $\lambda/(1 - \alpha)$ and thus decreasing in α . Thus, as ε increases, ρ becomes smaller.

Third, the same trick can be used to show that ρ is increasing in λ , while ε is decreasing in λ . QED ■

We now turn to formal tests of this proposition. We do not observe learning effort in the data, but we know from the model that learning effort is high for funds who still receive funding in state D , i.e. when past performance has been relatively poor. We use the fact that, in our data, funds with impediments to withdrawal face lower reductions in assets under management conditional on bad performance (see also Ding et al, 2008, for related evidence). Thus, we use the presence as strong impediments to withdrawal as our measure that $I_D > 0$, and test the two predictions derived here.

3 Empirical Section

3.1 Data Description

We start from a June 2008 download of EurekaHedge, a hedge fund data provider. The download provides us with monthly data from June 1987 until June 2008. 6,070 funds are initially present in the sample, with a total of 366,728 observations. Every month, each fund reports asset under management and net of fee returns. We delete from the data set all funds that have less than \$20m under management.

Our main results use annual data, but we also use higher frequency information on returns (monthly and quarterly, see below). Descriptive statistics on returns, and AUM are provided at the annual frequency in Table 1, panel A. Mean annual return about 11% net of fees. Mean assets under management are 300 million dollars. Also available from the data

are fund level characteristics that do not change over time, whose descriptive statistics are reported in panel B, Table 1. Using these information, we construct two dummy variables capture the presence of “strong” impediments to withdrawal:

- **Lock up dummy:** In some cases, investors agree to lock their investment in the fund for a given period of time for certain length of time after their investment.. Out of 5,154 funds for which share restrictions are known, the mean lock-up period is about 2.6 months. This mean conceals a lumpy distribution: 21% of the funds have lock up periods, 15% have a lock up period of 12 months, and only 2% have a longer lock up period. The percentage of funds with lock-up periods that we have in our dataset is similar to what Aragon (2007) has in his TASS extract.
- **Redemption dummy:** Once the lock-up period has passed, investors can redeem their shares, but still face constraints. Redemption can only occur at fixed moments of the year. For 3,152 funds (53% of the total), redemption is monthly. It is quarterly for 25% of the funds (1,499), and annual in 151 cases. In addition, investors have to notify the fund of their withdrawal before the redemption period. This notice period is lower than 1 month in 30% of the cases, equal to 1 month in 30% of the cases, and is equal or above one quarter in 15% of the cases. We construct a dummy variable equal to one when the sum of the redemption and notice periods is equal or longer than a quarter (90 days). The mean value of this sum is equal to 92 days; for 38% of the funds, it is equal or larger than a quarter.

Finally, the spearman correlation between the lock-up dummy and the redemption dummy is 41% (using one data point per fund). Thus, even though this correlation is positive and

statistically significant, which indicates some complementarity between the two forms of share restriction, it is far from being equal to 1. In particular, 23% of the funds without lock up have “redemption periods”.

[Table 1 about here]

How effectively constrained are hedge fund investors? To answer this question, we compute the mean duration of capital, for each fund, separately for each year. We do this by including the effects of lock up periods, redemption date and advance notice. We use the following formula:

$$\text{Duration}_{it} = \text{Notice}_i + \frac{\text{Redemption Period}_i}{2} + \frac{1}{\text{AUM}_{it}} \cdot \left(\sum_{0 \leq s \leq \text{L-U pd}_i} \text{Net Infl}_{it-s} \times 1_{\{\text{Net Infl}_{it-s} > 0\}} \times (\text{L-U Pd}_i - s) \right)$$

The first part of this formula accounts for the effect of notice and redemption periods. The implicit assumption behind this formula is that fund’s distance to the next redemption period is uniformly distributed. The second part accounts for the effect of lock up periods. For each past net inflow into the fund, it computes the remaining lock up duration (for instance, 5 month old inflows have a duration of 7 months if the lock up period is one year). We then normalize by current assets under management. We use monthly data. Following the literature on fund flows (Chevalier and Ellison, 1997, Sirri and Tufano, 1998), we compute net inflows by taking the difference between monthly AUM growth and monthly returns, and remove outliers. Overall, the above formula is an approximation. First, past inflows are computed net of outflows. This procedure leads us to underestimate gross inflows if they occur at the same time as gross outflows. Second, when shares are still locked up, the notice

and redemption periods are in part ineffective. This leads the above formula to overestimate duration.

[Figure 1 about here]

We plot the sample distribution of estimated durations in Figure 1. Taking all fund-months in the sample, the sample mean of this measure is 3 months. On average, the contributions of potential lock up periods and redemption and notice are of similar sizes. The 25th, 50th and 75th percentiles of the distribution are respectively 1, 1.5 and 3.5 months. The time series of the mean duration exhibits a clear downward trend, from 3.5 months in 1996 to about 2.6 months in 2007. If we focus on the subgroup of funds with lock up periods (21% of our sample), mean duration is, unsurprisingly, much larger: 8.2 months (median is 5.8). Thus even though most funds have relatively short duration of liabilities, there is a group of funds for which the mean dollar of AUM is secured for at least half a year.

As expected, such impediments to withdrawals do indeed prevent outflows from happening in the data. To check this, we run the following regression on annual data:

$$\text{Outflow}_{it} = \gamma_i + \beta \cdot 1_{\{r_{it-1} < r_{t-1}^{rf}\}} + \delta \cdot 1_{\{r_{it-1} < r_{t-1}^{rf}\}} \times \text{Impediment}_i + \varepsilon_{it}$$

where Outflow_{it} is a variable equal to 0 if the fund experiences net inflows in year t , and equal to net inflows if net inflows are negative. Net inflows are computed as the difference between AUM growth (between t and $t - 1$) and net-of-fee returns, as is standard in the literature. $1_{\{r_{it-1} < r_{t-1}^{rf}\}}$ is a dummy variable equal to 1 if past year's returns have been lower than the safe rate of return (as measured by the yield on 3 months Treasury bill). Impediment_i is one of the two measures described above: lock-up period of at least a year, or redemption period

of at least a quarter. We include a fund specific fixed effect γ_i and cluster error terms at the year level.

[Table 2 about here]

Regression results are reported in Table 1. As shown in the first column, if past performance is below the safe rate of return, outflows increase by 11% of AUM on average. This is sizeable, compared to mean annual outflows of 13% in the data, and a cross sectional standard deviation of 22% (see Table 2). As shown in columns 2 and 3, such a large sensitivity is somewhat reduced, yet not totally erased, by the presence of impediments to withdrawal. Conditional on low performance, funds with such share restrictions experience outflows of 8% of AUM, compared to 12% without such restrictions. Hence, the sensitivity is reduced by about one third.

3.2 Evidence from Conditional Returns

We test here our prediction 1: illiquid funds overperform liquid funds relatively more after bad performance than after good performance. We first provide graphical evidence in figure 2. To obtain it, we regress current and past returns on fund dummies to absorb individual fixed effects. Then, we plot the residuals of current returns against the residuals of past returns for locked up funds and for liquid funds. It appears from figure 2 that the returns of locked funds exhibit more mean reversion than the returns of liquid funds. It also appears that, consistently with the model, such mean reversion is more prevalent when past performance is low. In the following, we test these two dimensions (excess mean reversion, and asymmetry) in turn.

[Figure 2 about here]

We first run the following regression:

$$r_{it} = \gamma_i + \beta \cdot 1_{\{r_{it-1} < r_{t-1}^{rf}\}} + \delta \cdot 1_{\{r_{it-1} < r_{t-1}^{rf}\}} \times \text{Impediment}_i + \varepsilon_{it} \quad (10)$$

where r_{it} is the annual return of fund i in year t . We use annual data because at the annual frequency returns are less likely to be polluted by asset illiquidity problems (Lo, 2008; more on this below). γ_i is a fund-specific fixed effect, designed to capture heterogeneity in risk exposure and alphas, across funds (but our results are unchanged in the absence of fixed effects). We cluster error terms at the year level. Our theory predicts that the extent of mean reversion in returns should be larger for illiquid funds, i.e. $\delta > 0$.

[Table 3 about here]

Table 3 reports the results. Column 1 indicates that there is, in the data, a slight mean reversion of fund returns. After a performance below the risk-free rate, annual returns are on average 3.8 points higher, but the statistical significance is small. Columns 2 and 3 separate out liquid and illiquid funds. Consistently with the first prediction of our model, the mean reversion of returns significantly increases with impediments to withdrawal. After returns below the risk-free rate, annual returns increase by 2.6 points for funds with no lock-up and by 7.7 points for funds with a lock-up: the difference is strongly significant. When we look at redemption periods, we find that returns increases by 3.4 points after low performance when the notice + redemption period is shorter than a quarter, and by 6.7 when it is longer than a quarter. Again, the difference is strongly significant.

Thus, results from Table 3 show that there is a significantly larger tendency for returns of illiquid funds to mean revert, but it does not differentiate between mean reversion in bad states of nature and mean reversion in good states of nature. Our model, however, does predict such an asymmetry. Prediction 1 suggest that most of the mean reversion should be conditional on bad states of nature. This comes from the fact that, in the model, bad states of nature are states where assets are underpriced, while there is no mispricing in the good states of nature. Such a prediction holds even if $\mu = 1$, i.e. expected asset payoffs are similar in states U , M , and D .

To test it, we check if there is a difference in mean reversion between funds with past low returns, and events with past high returns. We define low returns as above, i.e. as cases when returns are below the safe rate of return (as measured by the 3 month.T-bill rate). This corresponds to approximately 21% of the observations in our dataset. Against this background, we define high returns as cases where past annual returns are above 20% net of fee: this threshold is chosen because it isolates about 21% of our fund-year observations, so the identifying power of the data should be similar for mean reversion conditional on high and low returns.

[Table 4 about here]

We then run regressions similar in spirit to (10): we also introduce a dummy variable equal to one if past returns have been above 20%, and its interaction with the fund's impediments to withdrawal. Results are reported in Table 4. In column 1, we do not control for impediments. It appears that the amount of mean reversion is fairly symmetric across the two extreme states, albeit weakly significant. Conditional on low past returns, expected returns are

larger by about 3 percentage points. Conditional on high past returns, expected returns are lower by about 2.7 ppnts. In column 2, we interact these two terms with the lock-up dummy. The asymmetry appears clearly: conditional on bad performance, funds with lock-ups overperform funds without lock-ups by about 5 percentage points. Conditional on good past performance, both categories of funds experience the same decline in expected returns. The same asymmetry is present with our second measure of impediment to withdrawal.

3.3 Mean Reversion, conditional on other funds Capital Structures

In this Section, we test our prediction 2. This prediction states that the difference between high and low effort funds expected returns, conditional on low past performance, is a decreasing function of α . When the fraction of funds with $I_D > 0$ decreases, the price difference between states D and M will expand, which will increase the scope for mean reversion for funds that invest in state D . In our model, α is determined in equilibrium but the proof of prediction 2 ensures that this negative covariation should be there as underlying parameters change.

Econometrically, prediction 2 rewrites, for fund i at date t operating on asset market s :

$$\begin{aligned}
 r_{ist} = & \gamma_i + \beta \cdot 1_{\{r_{ist-1} < r_{t-1}^{rf}\}} + \delta \cdot 1_{\{r_{ist-1} < r_{t-1}^{rf}\}} \times \text{Impediment}_i \\
 & + \zeta \cdot 1_{\{r_{ist-1} < r_{t-1}^{rf}\}} \times \alpha_s + \eta \cdot 1_{\{r_{ist-1} < r_{t-1}^{rf}\}} \times \text{Impediment}_i \times \alpha_s + \varepsilon_{it} \quad (11)
 \end{aligned}$$

where α_s corresponds to α in the model; α_s is the fraction of other illiquid funds. Ideally, to compute α_s , we should focus on pure players on the same asset market where fund i operates. It is, however, not feasible, first because in general funds may operate on several

assets markets, and also because the data does not provide us with the positions of funds on each market. Thus, we use the style classification of the database, implicitly assuming that funds operating within a given style buy the same assets. Thus, α_s will be the fraction, in the style where fund i is operating, of funds a lock-up period of at least 1 year, or the fraction of funds with at least a quarterly redemption. Our model unambiguously predicts that the difference in mean reversion should be lower in styles where there are more other funds with impediments to withdrawal, so that $\eta < 0$.

[Table 5 about here]

Table 5 reports estimates of equation (11) on our sample. Column 1 measures impediments as lock-ups. Conditional on past low performance, illiquid funds outperform by about 13 percentage points if there no other illiquid funds at all in the style. If the fraction of other funds with lock-ups is 20% (the sample average), the outperformance of illiquid funds goes down to 7 percentage points. The differential effect of α_s is statistically significant at 5%. In column 2, we use quarterly redemptions as our measure of impediments to withdrawals, both at the fund level and to compute α_s . For this specification, the result is less favorable to our theory: the estimate of η is insignificant and economically negligible.

3.4 Comparison with Existing Literature

Our main empirical result is thus that funds experience more mean reversion in returns when they have share restrictions. In this Section, we show how our results complement the existing literature on hedge funds.

Getmansky, Lo and Makarov (2004) have designed a measure of returns “smoothing” by

funds (named θ_0). This measure is econometrically complex to put in place, but the principle is to look at autocorrelation of monthly returns. If monthly returns are very autocorrelated, then it is likely that funds smooth returns across months to minimize volatility. Such a strategy is easier to put in place for assets whose prices cannot easily be marked to market, so θ_0 is also considered as a proxy for asset illiquidity. Consistent with the idea that impediments to withdrawals help funds to buy illiquid assets, such interpretation, Aragon (2007), Ding et al (2008) and Liang and Park (2008) have shown that high θ_0 funds also tend to have share restrictions.

This contradiction with our results (we find *less* returns persistence for funds with share restrictions) is only apparent, because we focus on annual returns, which are less likely to be smoothed. Illiquidity and window dressing issues should therefore generate less autocorrelation at this frequency. In fact, as we show in Tables 3 and 4, column 1, annual returns are more likely to mean-revert, for the average hedge fund in our sample.

This difference between existing results and our paper does not come from the fact that we are using a different dataset (most papers use Lipper/TASS), but really from the fact that we work with annual data. To check this, we look at the correlation between impediments to withdrawal and the autocorrelation of monthly returns. First, we compute, for each fund, a measure of the first order autocorrelation in monthly returns: we find an average of 0.10. This figure is consistent with what can be found in papers using other datasets: for instance, Lo (2008) finds that the mean first order autocorrelation of fund returns is 0.08 using the TASS/Lipper database (his Table 2.6). Consistently with Liang and Park (2008), we also find that autocorrelation is positively correlated with the presence of a lock-up period: in our dataset, the correlation is 0.08 (compared to 0.09 in their study),

statistically significant at 1%. Thus, our data generates the same pattern as existing papers: impediments to withdrawal are associated with more autocorrelated monthly returns.

To further check this, we ran regression (10) with *monthly*, instead of yearly, data, on all funds. We find, in the first column of Panel A of Table 6, that coefficient on low past performance is -0.38 and the interaction term is -0.14 , but significant at 10% only. Thus, at very high frequency, our data, like others, generate the positive autocorrelation pattern found in the literature.

[Table 6 about here]

This analysis suggests that the correlation between impediments to withdrawals and autocorrelation of returns is affected by two opposite forces: at the annual frequency, illiquid funds mean revert more because they take advantage of temporary mispricing, while at the monthly frequency, they exhibit more autocorrelation because they hold illiquid assets whose prices display a significant inertia. To disentangle these two effects, we run regression (10) on different groups of hedge funds, depending on the liquidity of the asset market they operate on. If we focus on funds managing liquid assets such as (long-short) equity funds, we find some evidence that persistence of monthly returns is weaker for funds with lock-up periods (Panel A, column 2), although the relation is not statistically significant. So even with monthly data, the evidence from “liquid styles” is more in line with our theory. For fixed income funds, the pattern is reversed (Panel A, column 3): for this strategy, as the existing literature would predict, share restrictions means more smoothing, and therefore more autocorrelation. In unreported regressions, we find similar results when we measure impediments to withdrawal with redemption periods.

With quarterly data, our predictions start to have more bite. In Panel B of Table 6, we re-estimate equation (10), using as the LHS variable quarterly, instead of monthly, returns. On the whole panel of funds, we find evidence of more mean reversion for funds with lock-up periods (Panel B, column 1), while there was less reversion at the monthly frequency (Panel A, column 1). For equity funds, the effect is even more spectacular (Panel B, column 2). For fixed income funds, the movement of persistence is still present at the quarterly frequency, but less significant (Panel B, column 3). In unreported regressions, we find similar results when we measure impediments to withdrawal with redemption periods.

4 Conclusion

In this paper, we have developed and tested a model of delegated fund management in equilibrium. The starting point was Shleifer and Vishny (1997): arbitrageur invest in a common market. Arbitrage opportunities may generate temporary underperformance. Our model explicitly models the contract that ties the investor and the fund manager. Because in some cases the underpricing can be so severe, we find that it is always optimal for the investor to commit not to redeem his shares, for some funds only. Hence, we predict that funds with share restrictions will outperform those without such restrictions after past bad performance. We find evidence consistent with this in the data.

5 References

Acharya and Viswanathan, 2007, “Moral Hazard, Collateral, and Liquidity”, *mimeo LBS*

Agarwal, Daniel and Naik, 2008, “Flows, Performance, and Managerial Incentives in Hedge Funds”, *forthcoming Journal of Finance*

Agarwal, Fung, Loon and Naik, 2008, “Risk and Return in Convertible Arbitrage: Evidence From the Convertible Bond Market”, *mimeo LBS*

Aghion and Bolton, 1992, “An Incomplete Contracts Approach to Financial Contracting”, *The Review of Economic Studies*

Ang and Bollen, 2008, “Locked Up by a Lock Up: Valuing Liquidity as a Real Option”, *mimeo*

Aragon, 2007, “Share Restrictions and Asset Pricing: Evidence From the Hedge Fund Industry”, *Journal of Financial Economics*

Brunnermeier and Pedersen, 2008, “Market Liquidity and Funding Liquidity”, *The Review of Financial Studies*

Campbell and Viceira, 2002, “Strategic Asset Allocation for Long Term Investors”, *Oxford University Press*

Casamatta and Pouget, 2008, “Fund Managers’ Contracts and Short Termism”, *mimeo Toulouse School of Economics*

Chen, Hanson, Hong, and Stein, 2008, “Do Hedge Funds Profit From Mutual Funds Distress?”, *mimeo Harvard University*

Chevalier and Ellison, 1997, “Risk Taking by Mutual Funds as a Response to Incentives”, *The Journal of Political Economy*

Coval and Stafford, 2008, “Asset Fire Sales (and Purchases) in Equity Markets”, *Journal of Financial Economics*

Ding, Getmansky, Liang and Wermers, 2008, “Investor Flows and Share Restrictions in

the Hedge Fund Industry”, *mimeo*

Getmansky, Lo and Makarov, 2004, "An Econometric Model of Serial Correlation and Illiquidity of Hedge Fund Returns", *The Journal of Financial Economics*

Gromb and Vayanos, 2002, “Equilibrium and Welfare in Markets with Financially Constrained Arbitrageurs”, *Journal of Financial Economics*”

Hombert, 2007, “Fire Sales: Revisiting the Market Equilibrium Approach”, *mimeo Toulouse School of Economics*

Innes, 1990, “Limited Liability and Incentive Contracting with Ex Ante Action Choices”, *Journal of Economic Theory*

Lerner and Schoar, 2004, “The Illiquidity Puzzle: Theory and Evidence From the Private Equity Industry”, *The Journal of Financial Economics*

Liang and Park, 2008, “Share Restrictions, Liquidity Premium, and Offshore Hedge Funds”, *mimeo*

Lo, 2008, “Hedge Funds: An Analytic Perspective”, *Princeton University Press*

Ramadorai, 2008, “The Secondary Market for Hedge Funds and the Closed-Hedge Fund Premium”, *mimeo*

Rampini and Viswanathan, 2008, “Collateral, Financial Intermediation, and the Distribution of Debt Capacity”, *mimeo Duke University*

Shleifer and Vishny, 1997, “Limits to Arbitrage”, *The Journal of Finance*

Sirri and Tufano, 1998, “Costly Search and Mutual Funds Flows”, *The Journal of Finance*

Stein, 2005, “Why Are Most Funds Open-End? Competition and the Limits of Arbitrage”, *The Quarterly Journal of Economics*

Stein, 2009, “Sophisticated Investors and Market Efficiency”, forthcoming, *The Journal*

of Finance

Tirole, 2006, “The Theory of Corporate Finance”, *Princeton University Press*

A Figures




Figure 1: Duration of Fund Liabilities

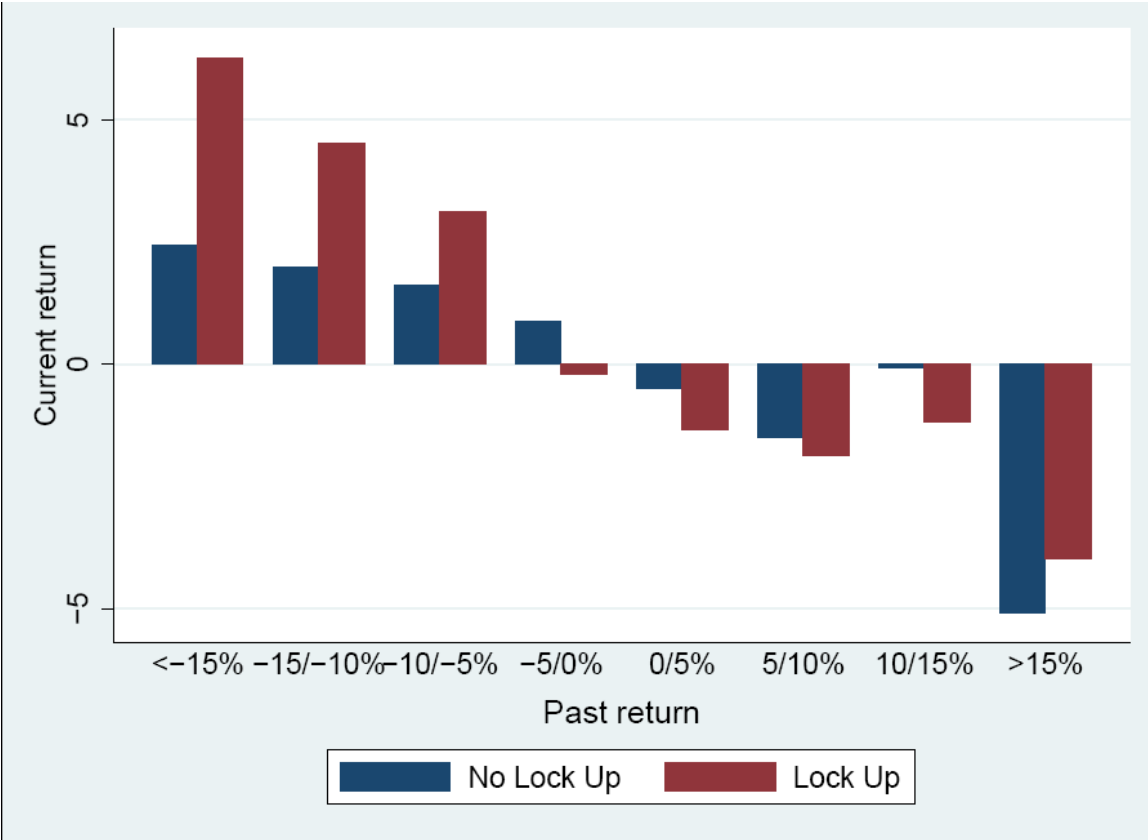


Figure 2: Conditional Returns and Lock-Ups

B Tables

	All funds with AUM>20m		
	Mean	Std	Obs
Panel A: Annual variables			
Return	10.5	13.5	7,041
AUM	301	554	7,041
Net Inflows	18.0	75.3	7,041
Net Inflows× (Net Inflows<0)	-13.0	22.6	7,041
Panel B: Fixed Characteristics			
Long Short	0.44	-	5,310
Global Macro	0.06	-	5,310
Fixed income	0.05	-	5,310
Lock-up Period (months)	2.6	5.8	5,154
Lock-up dummy	0.21	-	5,154
Notice + Redemption Periods	92	84	4,805
Quarterly Not.+Red. dummy	0.38	-	4,805

Source: Annual data from EurekaHedge, restricted to all funds with more than \$20m under management.

Table 1: Summary Statistics

Dependent variable	Net inflows _{it} × (Net inflows _{it} < 0)		
	Impediment to withdrawal	None	Lock Up
	(1)	(2)	(3)
$(r_{it-1} < r_{t-1}^{rf})$	-0.11***	-0.12***	-0.14***
	(0.2)	(0.2)	(0.02)
$(r_{it-1} < r_{t-1}^{rf})$	-	0.04**	0.04***
× Impediment _i		(0.02)	(0.01)
Fund FE	Yes	Yes	Yes
Observations	4,825	4,690	4,162
<i>Adj. R</i> ²	0.54	0.54	0.55

Data: EurekaHedge, 1994-2007. Annual data, excluding funds with AUM lower than 20 million USD. The dependent variable is equal to annual net inflows if they are negative, and zero else. Net inflows are computed as the difference between the growth in AUM minus net-of-fee returns. All specifications include fund specific fixed effects. In column (1), the only regressor is a dummy equal to 1 if the past annual return was lower than the yield on the 3 month T-bill. In column (2), we interact with the fact that fund i has a lock-up period of at least a year. In column (3) we interact with the fact that redemption + notice periods is at least 120 days. Error terms are clustered at the year level. *, **, and *** means statistically different from zero at 10, 5 and 1% levels of significance.

Table 2: Outflows and Impediments to Withdrawal

Dependent variable	r_{it}		
	None	Lock Up	Quart. Redemption
	(1)	(2)	(3)
$(r_{it-1} < r_{t-1}^{rf})$	3.8*	2.6	3.4
	(2.1)	(2.3)	(2.0)
$(r_{it-1} < r_{t-1}^{rf})$	-	5.1***	3.3***
$\times \text{Impediment}_i$		(1.4)	(1.0)
Fund FE	Yes	Yes	Yes
Observations	4,541	4,412	3,902
$Adj. R^2$	0.48	0.48	0.49

Data: EurekaHedge, 1994-2007. Annual data, excluding funds with AUM lower than 20 million USD. The dependent variable is the annual net-of-fee return. All specifications include fund specific fixed effects. In column (1), the only regressor is a dummy equal to 1 if the past annual return was lower than the yield on the 3 month T-bill. In column (2), we interact with the fact that fund i has a lock-up period of at least a year. In column (3) we interact with the fact that redemption + notice periods is at least 120 days. Error terms are clustered at the year level. *, **, and *** means statistically different from zero at 10, 5 and 1% levels of significance.

Table 3: Conditional Returns and Impediments to Withdrawal

Dependent variable	r_{it}		
	None	Lock Up	Quart. Red.
	(1)	(2)	(3)
$(r_{it-1} < r_{t-1}^{rf})$	3.0*	1.8	2.8
	(1.7)	(1.9)	(1.7)
$(r_{it-1} > 20\%)$	-2.7*	-2.9*	-2.1
	(1.5)	(1.5)	(2.0)
$(r_{it-1} < r_{t-1}^{rf})$ × Impediment _{<i>i</i>}	-	5.0***	2.8***
		(1.4)	(0.8)
$(r_{it-1} > 20\%)$ × Impediment _{<i>i</i>}	-	0.5	-1.0
		(1.0)	(2.1)
Fund FE	Yes	Yes	Yes
Observations	4,541	4,412	3,902
<i>Adj. R</i> ²	0.48	0.48	0.49

Data: EurekaHedge, 1994-2007. Annual data, excluding funds with AUM lower than 20 million USD. The dependent variable is the annual net-of-fee return. All specifications include fund specific fixed effects. In column (1), the two regressors are a dummy equal to 1 if the past annual return was lower than the yield on the 3 month T-bill, and a dummy equal to 1 if the past annual return was above 20%. In column (2), we interact with the fact that fund i has a lock-up period of at least a year. In column (3) we interact with the fact that redemption + notice periods is at least 120 days. Error terms are clustered at the year level. *, **, and *** means statistically different from zero at 10, 5 and 1% levels of significance.

Table 4: Conditional Returns and Impediments to Withdrawal: Good States versus Bad States

Dependent variable	r_{it}	
	Lock-Up	Quart. redemption
Impediment to withdrawal	(1)	(2)
$(r_{it-1} < r_{t-1}^{rf})$	0.1	4.5**
	(2.4)	(2.0)
$(r_{it-1} < r_{t-1}^{rf})$	12.7***	3.8
$\times \text{Impediment}_i$	(3.3)	(5.5)
$(r_{it-1} < r_{t-1}^{rf})$	13.7**	-4.4
$\times \alpha_{st}$	(5.1)	(3.3)
$(r_{it-1} < r_{t-1}^{rf})$	-31.1**	0.3
$\times \text{Impediment}_i \times \alpha_{st}$	(13.8)	(12.0)
Fund FE	Yes	Yes
Observations	4,412	3,902
$Adj. R^2$	0.49	0.49

Data: EurekaHedge, 1994-2007. Annual data, excluding funds with AUM lower than 20 million USD. The dependent variable is the annual net-of-fee return. All specifications include fund specific fixed effects. In column (1), we use as measure of impediment to withdrawal the fact that fund i has a lock-up period of at least a year. In column (2) the impediment dummy is equal to 1 if redemption + notice periods is at least 120 days. Error terms are clustered at the year level. *, **, and *** means statistically different from zero at 10, 5 and 1% levels of significance.

Table 5: Mean Reversion: The Impact of Other Funds' Illiquidity

Dep. Variable	r_{it}		
Panel A: Monthly frequency			
	All	Long short equity	Fixed Income
	(1)	(2)	(3)
$(r_{it-1} < r_{t-1}^{rf})$	-0.38**	-0.44**	-0.45***
	(0.18)	(0.21)	(0.14)
$(r_{it-1} < r_{t-1}^{rf})$	-0.14*	0.14	-0.49***
× Lock-Up _i	(0.08)	(0.13)	(0.14)
Fund FE	Yes	Yes	Yes
Observations	120,734	51,963	6,929
<i>Adj. R</i> ²	0.06	0.05	0.10
Panel B: Quarterly frequency			
	All	Long short equity	Fixed Income
	(1)	(2)	(3)
$(r_{it-1} < r_{t-1}^{rf})$	-0.10	-0.42	-0.49
	(0.42)	(0.58)	(0.40)
$(r_{it-1} < r_{t-1}^{rf})$	0.53**	1.37***	-1.24**
× Lock-Up _i	(0.26)	(0.24)	(0.55)
Fund FE	Yes	Yes	Yes
Observations	34,447	14,828	1,989
<i>Adj. R</i> ²	0.15	0.15	0.18

Data: EurekaHedge, 1994-2007. Annual data, excluding funds with AUM lower than 20 million USD. The dependent variable is the net-of-fee return. Panel A uses monthly returns, while Panel B uses quarterly returns. All specifications include fund specific fixed effects. In column (1), we look at all funds. In column (2), we restrict the sample to the long-short equity style. In column (3), we restrict the sample to funds operating in the fixed income style. Error terms are clustered at the month (panel A) and quarterly (panel B) level. *, **, and *** means statistically different from zero at 10, 5 and 1% levels of significance.

Table 6: Conditional Returns and Impediments to Withdrawal: Higher Frequency Evidence