

WHY ARE MOST FUNDS OPEN-END? COMPETITION AND THE LIMITS OF ARBITRAGE*

JEREMY C. STEIN

The majority of asset-management intermediaries (e.g., mutual funds, hedge funds) are structured on an open-end basis, even though it appears that the open-end form can be a serious impediment to arbitrage. I argue that when funds compete to attract investors' dollars, the equilibrium degree of open-ending in an economy can be excessive from the point of view of these investors. One implication of the analysis is that, even absent short-sales constraints or other frictions, economically large mispricings can coexist with rational, competitive arbitrageurs who earn small excess returns.

I. INTRODUCTION

This paper is motivated by two stylized facts. First, the vast majority of professionally managed investment vehicles (e.g., mutual funds, hedge funds) are structured on an open-end, as opposed to a closed-end basis, making it possible for their clients to liquidate shares on demand. Second, both theory and evidence suggest that the open-end form imposes serious constraints on arbitrageurs, since they are exposed to the risk of withdrawals if they perform poorly in the short run.¹ This risk can make it dangerous for them to put on trades that are attractive in a long-run sense, but where convergence to fundamentals is unlikely to be either smooth or rapid—a point that has been made by Shleifer and Vishny [1997]. For example, open-end funds are unlikely to want to bet heavily against something like the dot-com bubble of the late 1990s.²

Given the obvious drawbacks, why is the open-end form so

* I am grateful to the National Science Foundation for financial support, and to Adam Szeidl and Ann Richards for research assistance. Thanks also to Lucian Bebchuk, Markus Brunnermeier, Edward Glaeser, Hong Liu, Stefan Nagel, Michael Schwarz, Andrei Shleifer, Wei Xiong, the referees, and seminar participants at Brandeis University, the University of Chicago, Emory University, the Federal Reserve Bank of New York, Harvard University, the Massachusetts Institute of Technology, Rice University, the University of Michigan, and Washington University for helpful comments.

1. See, e.g., Ippolito [1992], Chevalier and Ellison [1997], and Sirri and Tufano [1998] for evidence on the flow-performance relationship in the mutual-fund sector, and Agarwal, Daniel, and Naik [2004] for analogous evidence regarding hedge funds.

2. Brunnermeier and Nagel [2004] document the reluctance of hedge funds to buck the dot-com bubble, and Lamont and Stein [2004] find that investors as a whole actually scaled back their short positions as the NASDAQ index was exploding upwards from mid-1998 through March of 2000.

dominant? One answer, in the survival-of-the-fittest spirit of Fama and Jensen [1983], might be that open-ending is an optimal response to agency problems.³ If a fund is set up on a closed-end basis, dispersed investors have no recourse in the face of managerial misbehavior, and may see their entire investment slowly eaten away. In contrast, if the fund is open-end, they can liquidate their positions at the first sign of trouble, thereby avoiding large losses due to mismanagement or theft.

In this paper I take issue with the survival-of-the-fittest view. While maintaining the premise that agency considerations play a crucial role in the decision to open-end, I show that when there is also asymmetric information about managerial quality, the end result may be a degree of open-ending that is *socially excessive*.⁴ The basic argument is as follows. First, suppose that from an efficiency perspective, it would be better if all funds were closed-end. This amounts to assuming that the gains from being able to undertake longer-horizon trades in the closed-end form outweigh the potential losses that come from being unable to control agency-prone managers.

Unfortunately, in such a case it may be impossible to sustain the efficient outcome where all funds remain closed-end. Any one high-ability manager will have an incentive to deviate, and set his fund up with an open-end charter, because this will signal his quality and thereby allow him to get more assets under management. This business-stealing effect sets off a counterproductive race to be open-ended, with both high-quality and low-quality managers ultimately being forced to open-end in order not to lose their investors. In the resulting equilibrium, all funds are open-end, yet there is no net benefit in terms of the quality of the managerial pool, since no low-ability managers get screened out. For some parameter values, the all-closed-end scenario can be Pareto-dominant, yielding higher returns to both investors and fund managers. But unlike the all-open-end configuration, it is not a stable equilibrium outcome.

3. Fama and Jensen [1983, p. 327] write: “. . . there is competition among organizational forms for survival in any activity. Absent fiat, the form of organization that survives in an activity is the one that delivers the product demanded by customers at the lowest price while covering costs . . . An important factor in the survival of organizational forms is the control of agency problems.”

4. Bebchuk [2002] makes an analogous point in a different context: he demonstrates that when there is asymmetric information at the IPO stage, nonfinancial firms going public may adopt corporate charters that do not maximize aggregate value.

This line of reasoning has implications for the debate over market efficiency. Proponents of the efficient-markets hypothesis often point to the relatively low net-of-fees excess returns earned by investors in professionally managed funds as Exhibit A in making their case. Their implicit argument has two parts: first, that minimal barriers to entry and intense competition in the money-management industry have driven aggregate arbitrage profits close to zero; and second, that small arbitrage profits must reflect an absence of economically significant mispricing. Ross [2001] writes:

But, if we allow for inefficiency, how much is there? As a rough guide, the hedge funds listed in the offshore directory through about 1996 had a (generously estimated) average alpha . . . of about 4 percent per year. With about \$200 billion of funds that means about \$8 billion of alpha per year. Assuming this undercounts banks and others engaged in similar activities by a factor of, say, 2, that brings the total to \$16 billion and doubling it again to account for all the missing other players takes us to \$32 billion which we will round up to \$40 billion just to be conservative. This is a tidy sum, but not when measured against a rough estimate of \$50 trillion of marketed assets. This means that, on average, prices are “off” or inefficient to about \$40 billion/\$50 trillion which is less than 0.1 percent. There should be no debate, then, over whether the efficiency glass is half full or half empty; it is simply quite full.

The point made here is that even if the money-management industry is highly competitive, and aggregate arbitrage profits are therefore small, it does not follow that markets are close to efficient. Indeed, competition among fund managers may be part of the problem, rather than part of the solution. For this competition can lead to a situation in which all funds are driven to an inefficient organizational form, namely being open-ended. Given this inefficient form, funds will stick primarily to short-horizon strategies, and earn low excess returns. In so doing, they will leave large long-horizon mispricings such as the dot-com bubble mostly untouched, because attacking such mispricings aggressively would require a closed-end structure.

This paper can be seen as complementary to Shleifer and Vishny [1997]. Shleifer and Vishny begin with the *assumption* that professional arbitrage is conducted by open-end funds, and then go on to explore the implications of the ensuing performance-flow relationship for asset prices. They show that, in extreme circumstances, professional arbitrage can have a destabilizing effect on prices: as asset prices diverge farther from fundamentals, arbitrageurs lose money, suffer outflows, and may have

to engage in “fire-sale” liquidations of their positions. More generally, a central message of their paper is that open-end arbitrageurs will tend to be timid in attacking mispricings in the first place, precisely because they want to avoid such forced-liquidation scenarios.

In contrast, this paper focuses on the question of *why* the open-end form is so prevalent in the money-management industry.⁵ By demonstrating that the extent of open-ending can be socially excessive, the analysis suggests that the limits to arbitrage associated with the open-end form may be greater than previously appreciated. To see why, imagine a critique of Shleifer and Vishny [1997] along the following lines:

Perhaps arbitrage by open-end funds does go wrong in certain extreme circumstances. But these circumstances must be quite rare—and relatively inconsequential from an ex ante perspective—or else there would be strong incentives to organize funds differently, i.e., on a closed-end basis. Since we don't see many closed-end funds, the impediments to arbitrage created by open-ending must be minimal.

The logic developed above suggests that this critique is off the mark: in a world where funds compete for investors' dollars, it is not generally correct to presume that they will be organized in such a way as to be maximally effective at the task of arbitrage.

This paper draws on the insights of a large literature in corporate finance on the costs and benefits of short-term debt.⁶ Diamond [1991], in particular, makes a point similar to that here: firms that have positive private information about their credit quality may issue short-term debt as an optimistic signal, even though this increases the risk that they will face inefficient early liquidation (see also Flannery [1986]). The general theme is that

5. Chordia [1996] and Nanda, Narayanan, and Warther [2000] study mutual funds' use of exit fees, which, like a closed-end structure, also deter redemptions. However, their models have a different focus than mine—they assume that managerial ability is known, and that investors may have interim liquidity needs. Thus, back-end fees are seen as an efficient way to screen out high-liquidity-demand investors, whose frequent redemptions impose a transaction-cost externality on more stable investors.

6. General references include Bolton and Scharfstein [1990, 1996] and Hart and Moore [1998]. In the banking literature, several papers ask why banks choose to issue demandable liabilities, when demandability can interfere with their ability to carry out their lending function. According to Calomiris and Kahn [1991], Flannery [1994], and Diamond and Rajan [2000, 2001], demandable claims are a way of reassuring investors in a bank that they will not be too badly expropriated by bank management.

with asymmetric information, there can be excessive “commitment aversion” in financial contracting, since bad types benefit more from locking in their investors, and nobody wants to be taken for a bad type. In the specific context analyzed here, the cost of commitment aversion is that it interferes with long-horizon arbitrage.

The remainder of the paper is organized as follows. In Section II, I present some evidence on the dominance of the open-end form in the mutual-fund and hedge-fund sectors. The formal model is developed in Section III. Section IV concludes.

II. OPEN-END VERSUS CLOSED-END FUNDS: BASIC FACTS

In the mutual-fund sector the distinction between an open-end fund and a closed-end fund is completely clear: investors in an open-end fund can liquidate their shares on demand at any time, while investors in a closed-end fund can never liquidate, absent a takeover of the fund or other terminating event. Table I summarizes the breakdown of U. S. mutual funds over the period 1996–2002. At the end of 2002 there were 7,267 open-end funds, with \$4.1 trillion of assets, across equity, bond, and hybrid categories. By contrast, there were only 562 closed-end funds, with \$156.4 billion of assets. Thus, open-end funds accounted for 96 percent of the sector’s assets, and 93 percent of funds. As the table shows, the pattern in other recent years has been essentially identical.

TABLE I
OPEN-END VERSUS CLOSED-END MUTUAL FUNDS, 1996–2002*

Year	Open-end funds**		Closed-end funds***		Open-end/total	
	Funds	Assets (\$B)	Funds	Assets (\$B)	Funds	Assets
1996	5,260	2,624.0	475	144.5	0.92	0.95
1997	5,671	3,409.3	482	148.9	0.92	0.96
1998	6,288	4,173.5	490	153.0	0.93	0.96
1999	6,746	5,233.2	503	142.8	0.93	0.97
2000	7,116	5,119.4	483	136.9	0.94	0.97
2001	7,292	4,689.6	528	139.7	0.93	0.97
2002	7,267	4,119.6	562	156.4	0.93	0.96

* Source: *Mutual Fund Fact Book* 2003, Investment Company Institute.

** Includes all open-end U. S.-based equity, hybrid and bond funds; excludes money market funds.

*** Includes all closed-end U. S.-based equity and bond funds.

The hedge-fund sector is substantially smaller than the mutual-fund sector: as of year-end 2002, there were 4,600 U. S. hedge funds, managing a total of \$340 billion across all investment categories.⁷ Moreover, the distinction between an open-end and a closed-end hedge fund is not as black-and-white. Virtually all hedge funds allow their investors to liquidate their positions at some horizon; in this sense, they are all quasi-open-end. At the same time, most hedge funds put some restrictions on withdrawals, and so do not represent as pure a case as open-end mutual funds.

According to a recent survey, the most common arrangement, used by 43.6 percent of U. S. hedge funds, is quarterly redemption; i.e., investors can liquidate at three-month intervals, typically after giving 30 days' notice of their intention to do so. Another 25.7 percent of funds allow only annual redemption, while 15.7 percent are more flexible, allowing monthly redemption. Less than 1 percent of U. S. hedge funds allow the kind of daily redemption that is typical in the mutual-fund sector. Also, redemptions are only possible after the expiration of an initial lockup period. In the survey, 58.0 percent of U. S. funds impose lockups, most often of one year in length.⁸

The model below is not rich enough to capture all these aspects of fund structure: funds only face a simple dichotomous choice of whether to be open-end or closed-end. However, it would seem that the basic idea—that competition can lead to an inefficient degree of commitment aversion—is more general, and should apply to other, more continuous dimensions of structure that can also act as limits to arbitrage. For example, in the hedge-fund sector, where pure open-ending is very rare, the model nevertheless suggests that redemption intervals and lockup periods may be shorter than optimal.

7. These figures are from the website of Van Hedge Fund Advisors International, *www.hedgefund.com*. Worldwide, there were 7,500 funds with \$650 billion in assets at year-end 2002.

8. Source: Cerulli Associates [1999]. Cerulli in turn seems to draw much of their primary data from Van Hedge Fund Advisors. Interestingly, the figures for hedge funds worldwide suggest a somewhat more liberal attitude toward redemptions than seen in the United States: 6.8 percent allow daily redemptions, 5.0 percent allow weekly redemptions, and 26.7 percent allow monthly redemptions.

III. THE MODEL

III.A. Assumptions

There are four dates: time 0, time 1, time 2, and time 3. At time 0, funds choose their charters, and decide whether to be open-end or closed-end. Investors place their money in the funds they expect to provide the highest returns; if there is a tie, the money is split equally among the most attractive funds. There are n funds and n dollars of capital to be allocated. Any fund has a capacity limit, and can handle up to a maximum of m dollars, where $1 < m < n$. Ex ante, each fund manager has an independent probability p of being a “good” type, and a probability $(1 - p)$ of being a “bad” type; the type is private information at time 0. All managers and investors are risk neutral.

After the funds are set up and the money has been allocated, each fund manager gets another private signal at time 1, which indicates whether or not he has any good trading ideas. In particular, a manager learns at time 1 that he is either “inspired” or “uninspired.” A manager who is of the bad type at time 0 learns that he is uninspired with certainty at time 1. A manager who is of the good type at time 0 has a probability q of being inspired at time 1, and a probability $(1 - q)$ of being uninspired. Thus, from an unconditional ex ante perspective, the probability of a manager being inspired is pq .

Fund managers also choose their investment strategies at time 1. An inspired manager can invest in either a “smooth-convergence” arbitrage trade, or an “uncertain-convergence” arbitrage trade. I take the payoffs on these trades to be exogenous, and refer to them as the SC and UC assets, respectively. The assets are perfectly divisible, so an inspired manager can choose a portfolio that is any linear combination of the two. The choice of which trades to undertake is noncontractible, so a fund cannot promise its investors that it will stick to a particular investment strategy.

An uninspired manager has no ability to spot arbitrage opportunities, and simply invests in a wasting asset that gradually loses money over time. All trades are put on at time 1, and can either be kept in place until time 3, or interrupted at time 2. If a fund’s charter specifies it to be open-end, investors have the right to liquidate its positions at time 2, if they so desire. If a fund is closed-end, investors do not have this right.

Any trading strategy yields two types of payoffs: verifiable returns that can be appropriated by investors; and nonverifiable private benefits that accrue to the fund manager. In the current context the private benefits are perhaps best thought of as kickbacks that managers get from directing trades to particular brokers. In keeping with this interpretation, I assume that a manager's total private benefits are proportional both to the volume of assets he has under management, as well as to the duration of his tenure.

The private benefits are a modeling device that lead managers to care about both attracting and retaining investors. One would think that the same thing could be accomplished by assuming—realistically—that managers charge investors an explicit fee that is proportional to assets under management. However, if one assumes perfect competition among funds (along with zero marginal costs), any such nonperformance-based fees are necessarily driven to zero, which eliminates the fund-size motive for managers. By contrast, since private benefits cannot be alienated, fund managers cannot compete with each other by lowering the private benefits they take.

For any trading strategy, both investment returns and private benefits depend on whether the position is liquidated early, at time 2, or held until time 3. The following table details the possible outcomes. In each case, an ordered pair $\{x, y\}$ denotes the associated gross investment returns and private benefits *per dollar invested*, respectively. The underlying discount rate is zero, so these can all be thought of as discounted payoffs.

Payoffs Under Different Scenarios, Per Dollar Invested

{Gross Return to Investors, Private Benefit to Manager}

Uninspired Manager

Hold until $t = 3$ $\{K, Tz\}$
 Interrupt at $t = 2$ $\{L, z\}$

Inspired Manager, Smooth-Convergence Trade

Hold until $t = 3$ $\{R^{SC}, Tz\}$
 Interrupt at $t = 2$ $\{1, z\}$

Inspired Manager, Uncertain-Convergence Trade

Hold until $t = 3$ $\{R^{UC}, Tz(1 + \alpha)\}$
 Interrupt at $t = 2$ $\{r, z\}$; where $r = x$ with prob γ ,
 and $r = 1$ with prob $(1 - \gamma)$.

The key parametric assumptions with respect to investors' returns are as follows. First, $K < L < 1$, which implies that an

uninspired manager dissipates investors' money slowly but surely; some is gone by time 2, and even more is lost by time 3. Second, $R^{UC} > R^{SC} > 1$, which means that if an inspired manager can hold a position until time 3, returns are certain to be higher if he invests in the UC asset. Third, r is a random variable, whose realization is independent across managers, and is given by $r = x$ with probability γ , and $r = 1$ with probability $(1 - \gamma)$, where $x < L$. Thus, while the UC asset always does better than the SC asset if it is held until time 3, it is riskier from the perspective of early liquidation; this is the idea of uncertain convergence. In particular, there is a probability γ that if the UC asset is sold at time 2, investors will realize a loss on the position—a loss greater than they would suffer with an uninspired manager.

With respect to private benefits, there are two assumptions. First, $T > 1$, so that managers' payoffs are reduced with early liquidation. Second, $\alpha > 0$, which represents a form of incentive alignment—if he is sure not to be interrupted before time 3, an inspired manager prefers to invest in the UC asset, as his investors would like him to. This incentive alignment might reflect the fact that there is more scope for a fund manager to generate kickbacks on trades if the assets he is managing have appreciated substantially in value. Alternatively, if one is willing to step a bit outside of the model, the α parameter could be thought of as a proxy for the extent to which the fund manager is given a share of long-run returns; this would fit with the standard practice in the hedge-fund industry.⁹

III.B. Analysis

Let us start by examining what happens if all funds are closed-end. First, observe that Bertrand competition among funds with zero marginal costs drives fees to zero, since the total capacity of the fund sector mn , exceeds the amount of capital n available for investment. So the only payoffs to managers are the private benefits they extract. If investors cannot distinguish the good managers from the bad, every fund gets one dollar of assets to manage. Given that they face no danger of early liquidation in the closed-end form, inspired managers invest in the UC asset, and investors' gross expected return is $pqR^{UC} + (1 - pq)K$. Each inspired manager gets $Tz(1 + \alpha)$, and each uninspired

9. Explicit incentive contracts of the sort seen in the hedge-fund sector are discussed below.

manager gets Tz . As long as $pqR^{UC} + (1 - pq)K > 1$ (which I assume henceforth), investors earn a positive net return, so their participation constraint is satisfied. If p and q are sufficiently large, this closed-end scenario can represent a relatively efficient outcome. There are losses when managers turn out to be uninspired, but these are offset by the high returns that accrue when inspired managers invest in the UC asset.

The all-closed-end situation described above is a perfect Bayesian equilibrium. But it may not be robust, in the sense that it may rest on out-of-equilibrium beliefs that are “unreasonable” and hence that do not survive standard refinements, such as that of Cho and Kreps [1987].¹⁰ To investigate this issue, let us entertain the possibility that a good type deviates from the all-closed-end pooling equilibrium by i) switching to an open-end form; and ii) announcing that he will continue to not charge any fees.

In this context, a Cho-Kreps [1987] type of argument will deem the all-closed-end equilibrium to be nonrobust under the following conditions.¹¹ First, a bad type must not want to make the specified deviation (to being open-end with no fees) even if this would lead him to be judged a good type. Second, a good type must prefer to make the deviation if this leads him to be judged a good type. Implicit in the second condition is a third—namely that investors prefer to give their money to a deviator if they think he is a good type, even though he is now open-end, which may adversely affect his returns.

Intuitively, if these conditions are satisfied, a good type will break the all-closed-end equilibrium by setting up as an open-end fund and implicitly saying to investors: “Give your money to me—my willingness to go open-end proves that I’m a good type. For if I were a bad type, I would never want to go open-end, no matter what I could make you believe about me by doing so.” This reasoning leads to the following proposition.

10. In particular, the all-closed-end pooling equilibrium can always be supported by the out-of-equilibrium belief that a deviator who open-ends is a bad type, and hence should not be given any money to manage. The question is whether this belief is a plausible one.

11. Strictly speaking, the Cho-Kreps [1987] criterion is for static games. I am extending the criterion to the current dynamic setting by saying that out-of-equilibrium beliefs at the first stage of a game are unreasonable if there are no continuation equilibria following a deviation where the type described in the beliefs (here the bad type) can be made better off. An equilibrium action profile is then deemed to be not robust if it cannot be supported by reasonable beliefs.

PROPOSITION 1. If the following four sufficient conditions hold, the all-closed-end pooling equilibrium is not robust to a Cho-Kreps [1987] type of refinement; i.e., it cannot be supported by “reasonable” out-of-equilibrium beliefs.

$$(1) \quad (R^{UC} - x)q\gamma < (L - K)(1 - q),$$

$$(2) \quad m < T,$$

$$(3) \quad qmTz + (1 - q)mz > qTz(1 + \alpha) + (1 - q)Tz,$$

$$(4) \quad qR^{SC} + (1 - q)L > pqR^{UC} + (1 - pq)K.$$

The proposition is formally proved in the appendix, but a sketch of the argument goes as follows. First, condition (1) ensures that rational investors always liquidate any fund whose time-2 performance is less than or equal to L . That is, given their time-0 prior (stipulated in the out-of-equilibrium thought experiment) that the manager is certainly a good type, their time-2 posterior beliefs, conditional on a poor performance of L or less, attach a sufficiently high probability to the manager being uninspired that liquidation has a higher expected payoff than continuation.¹²

Given (1), a bad type knows that he will be liquidated for sure at time 2 if he switches to open-end, since he is certain to be uninspired and hence have a poor interim performance of L . So even if the deviation initially convinces investors that he is the best fund manager available, and lands him m dollars to manage, his total private benefits will only be mz . So long as these are less than Tz , the private benefits he gets in the all-closed-end equilibrium, he will not want to deviate—hence condition (2).

Condition (3) ensures that a good type *will* prefer to deviate to open-end if this lands him m dollars to manage. The way it is stated, (3) is actually more restrictive than it needs to be, since it implicitly assumes that if the good type turns out to be inspired, he will be excessively conservative in the open-end form, invest-

12. An implicit assumption is that investors' beliefs about managerial ability are such that a manager whose time-2 portfolio value is less than L cannot be judged to be better than a manager with performance equal to L . For example, a manager whose portfolio has a time-2 value of $x < L$ cannot be judged to be inspired, even though uninspired managers actually never do worse than L . One way to rationalize this assumption is to imagine that an uninspired manager always has the option to *underreport* (but not *overreport*) his time-2 performance. In this case, it can never be incentive-compatible for a performance of less than L to lead to a more positive assessment than a performance of L . I discuss this point in greater detail below.

ing everything in the SC asset. As will be discussed in more detail shortly, an inspired manager can do strictly better than this—even if he wants to keep the probability of liquidation at zero, he can still afford to put some (possibly small) fraction of his portfolio in the UC asset, which would increase his private benefits. At the expense of a little extra notational complexity, this logic leads to a weaker version of (3), which is described in the appendix.

Finally, condition (4) is sufficient for investors to prefer an open-end fund that they know to be managed by a good type, as opposed to a closed-end fund with a manager of unknown type. As with (3), condition (4) is also more restrictive than it needs to be. It again presumes that an inspired manager in the open-end form puts everything in the SC asset, when in fact he can always do strictly better than this.

By inspecting the conditions, we can get a sense of the circumstances under which closed-end funds might potentially survive. First, (4) suggests that the all-closed-end pooling equilibrium is more likely to be viable when p is near one, so that there is relatively little chance of the manager being a bad type.¹³ For example, somebody like Warren Buffett, with a very strong reputation, effectively has a p close to one. This is why he can use Berkshire-Hathaway—a closed-end corporation—as his investment vehicle, and still not have trouble raising money. Relatedly, Agarwal, Daniel, and Naik [2004] find that hedge funds that carry more restrictions on withdrawals (longer lockups, less frequent redemptions) have better performance. One interpretation is that only managers with high *ex ante* ability assessments are able to impose such restrictions in equilibrium.

Second, according to both (1) and (4), closed-end funds are more likely to be observed in situations where the time-2 liquidation value L is relatively low. In such a case, early interruption is so costly that the prospect of it can make deviation to the open-end form unattractive. This may explain why, e.g., emerging-market country funds—which tend to deal in relatively illiquid securities—are more often closed-end than other types of equity funds. In a similar vein, it is also not surprising that

13. To map the model into real-world data, think of there being multiple pools of fund managers, with each pool having a different (publicly observed) value of the ability prior p . Thus, even if there is a unique equilibrium outcome for all managers in a given pool—e.g., high- p managers remain closed-end, while low- p managers go open-end—there will be different outcomes across the population.

venture-capital funds are effectively closed-end, since their assets are stakes in untraded private firms.

Irrespective of whether the conditions in Proposition 1 hold, it is straightforward to establish that the model never admits a separating equilibrium in which good and bad types reveal themselves through their time-0 choice of form.

PROPOSITION 2. There can never be a separating equilibrium in which good funds are open-end, while bad funds are closed-end. Nor can there be a separating equilibrium in which good funds are closed-end, while bad funds are open-end.

To prove the proposition, note that in any separating equilibrium, a bad fund is recognized as bad. This implies that it gets no money to manage, whether it is open-end or closed-end, because in either case, investors are unable to recoup their initial investment. In contrast, if a bad fund mimics the strategy of a good fund, it will expect to get a positive amount (approximately $1/p$ dollars) to manage, since in a separating equilibrium there are on average pn good funds managing a total of n dollars. Thus, for any separating equilibrium that is proposed, a bad type will always prefer to mimic a good type; i.e., a bad type's incentive-compatibility constraint can never be satisfied.

However, it is interesting to note that if one is willing to sprinkle in a little investor irrationality, it may be possible to sustain a particular separating equilibrium—one in which good funds are open-end, and bad funds are closed-end. Suppose that a small fraction k of investors are naïve, and randomly allocate their money across all funds, irrespective of whether they are open-end or closed-end. The remaining investors are rational, and channel all their money to the good (i.e., open-end) funds. Thus, in equilibrium each bad fund gets k dollars to manage, and each good fund gets $k + (1 - k)/p$. A bad fund will no longer deviate from the proposed equilibrium if $kTz > (1 - k + pk)z/p$. In other words, if a bad fund can still get a few naïve investors to give it money in the closed-end form—and can take each of them for a two-period ride that yields a large private benefit of Tz —it may make sense for it to stay closed-end.

Although I will leave it aside in what follows, and continue to focus on the case where all investors are rational, this naïve-investor variation may help to add empirical realism to the model. Without it, the model predicts that closed-end vehicles are always run by better-than-average managers; this is the Warren

Buffett effect mentioned above. However, this implication is at odds with the large literature on discounts among closed-end mutual funds, which is hardly suggestive of uniformly superior ability on the part of closed-end fund managers.¹⁴ Indeed, it seems more realistic to believe that closed-end funds inhabit both ends of the quality spectrum. In order to get the model to fit this stylized fact, it is necessary to add the naive investors.

If we continue to suppose that the conditions laid out in Proposition 1 all hold, we now know that not only is the all-closed-end pooling equilibrium ruled out, but so is any separating equilibrium. Thus, if we restrict attention to pure-strategy equilibria, all that remains is a pooling equilibrium where all funds are open-end. In the appendix I prove:

PROPOSITION 3. Under the conditions of Proposition 1, there exists a pooling equilibrium in which all funds are open-end. This equilibrium *is* robust to a Cho-Kreps [1987] type of refinement.

The equilibrium described in Proposition 3 can be supported by the off-path belief that a fund which deviates to the closed-end form is a bad type. As demonstrated in the appendix, this belief is a reasonable one, in the sense of Cho and Kreps [1987].

An interesting normative question to ask is whether investors and fund managers can be worse off in this all-open-end equilibrium than they would be in a world where all funds were closed-end. In other words, can it ever be possible to create a Pareto-improvement with a regulation that prohibits funds from open-ending?¹⁵

Recall that in the all-closed-end setting, investors have an expected return of $pqR^{UC} + (1 - pq)K$, inspired managers get $Tz(1 + \alpha)$, and uninspired managers get Tz . If all funds are open-end, there are two different scenarios that can obtain. In the first of these, which I call the “aggressive-investment” scenario,

14. Depending on the sample and time period, researchers typically find that closed-end mutual funds trade at average discounts of 10 to 20 percent of net asset value. A model with only rational investors also has trouble with another related fact about closed-end funds: that they manage to go public in the first place—typically at a premium to net asset value—even though they subsequently tend to trade at discounts. See Lee, Shleifer, and Thaler [1990] for a survey of work in this area.

15. I should stress that this is only a thought experiment designed to illustrate the internal logic of the model. It is obviously not meant as a serious policy suggestion.

inspired managers continue to put their entire portfolio in the UC asset. This implies that inspired managers have a time-2 performance of $r = x$ with probability γ , and of $r = 1$ with probability $(1 - \gamma)$.

What happens to an inspired manager whose performance is poor (i.e., who has $r = x$) at time 2? Consider things from the perspective of investors. Investors do not directly observe whether managers are inspired, and must make an inference based on interim performance. In other words, after observing performance at time 2, investors will assign an updated probability to the manager being inspired. I make the plausible assumption that this posterior probability assessment cannot be better for time-2 performance below L than it is for performance equal to L . That is, I rule out the possibility that a manager with a time-2 performance of x is judged more favorably than a manager whose time-2 performance is $L > x$. This is tantamount to assuming that an uninspired manager with a performance of L can (and will) always strategically underreport, and say that he only had a performance of x .¹⁶

Given this assumption, investors update as follows in the aggressive-investment scenario. First, if a fund has a time-2 performance of strictly greater than L , investors know with certainty that the manager is inspired. Second, if a fund has a time-2 performance of L or less, investors attach a probability $\pi < 1$ to the possibility that the manager is inspired. A simple Bayesian calculation for π yields

$$(5) \quad \pi = pq\gamma / (pq\gamma + (1 - pq)).$$

The smaller is γ , the lower is π . This makes sense—if inspired managers only rarely have poor interim performance, when such poor performance is observed, it will tend to be attributed to the fact that the manager is uninspired. Given (5), it is easy to prove:

LEMMA 1. If condition (1) holds, then in the aggressive-investment scenario, a fund whose performance at time 2 is less than or equal to L will be liquidated. Thus, an inspired manager who invests aggressively faces a probability of liquidation of γ .

16. Thus, the only observed performance reports less than or equal to L will be reports of x , and these reports will come from both uninspired managers and inspired managers who have had bad luck.

Lemma 1 indicates that, in an open-end world, inspired managers may run the risk of being prematurely liquidated if they put everything in the UC asset. Note that with perfect information, investors would never want to liquidate an inspired manager. But given that they do not directly observe whether managers are inspired, and have to make an inference based on observed interim performance, it will under some circumstances be optimal for investors to pull the plug after a poor time-2 showing. Loosely speaking, this happens if the posterior likelihood of the manager being inspired is sufficiently low, and the incremental amount salvaged by getting out of an uninspired fund early (i.e., $(L - K)$) is large enough. Condition (1) turns out to be sufficient for this to be the case.

We can now ask whether the aggressive-investment scenario can be consistent with equilibrium behavior on the part of fund managers. If condition (1) holds, an inspired manager who invests aggressively knows that he will be liquidated with probability γ . Thus, his expected private benefits are $(1 - \gamma)Tz(1 + \alpha) + \gamma z$. Alternatively, an inspired manager has the option of investing "cautiously." I define cautious investment as a policy of putting a fraction θ of the portfolio into the SC asset, and a fraction $(1 - \theta)$ into the UC asset, where θ is the smallest allocation to the SC asset that ensures that the manager will not be liquidated at time 2; i.e., that guarantees that his time-2 performance will exceed L . (Whenever performance exceeds L , investors can infer with certainty that the manager is inspired, and so will never liquidate.) Fudging the trivial openness problem, we have that

$$(6) \quad \theta = (L - x)/(1 - x).$$

If he invests cautiously, an inspired manager has expected private benefits of $\theta Tz + (1 - \theta)Tz(1 + \alpha)$. If these exceed the benefits to be had from investing aggressively, there cannot be an aggressive-investment equilibrium, since given the updating rules associated with such an equilibrium, an inspired manager will always prefer to deviate to the cautious-investment strategy. This establishes:

PROPOSITION 4. Suppose that the conditions of Proposition 1 hold, so that the only robust equilibrium is one in which all funds are open-end. Inspired managers will necessarily invest cau-

tiously, with a fraction θ of their portfolios devoted to the SC asset, if the following condition is also satisfied:¹⁷

$$(7) \quad (1 - \gamma)Tz(1 + \alpha) + \gamma z < \theta Tz + (1 - \theta)Tz(1 + \alpha).$$

We have thus far established a set of conditions which imply both that funds will be open-end, and that this open-end form will in turn affect investment. It remains to compare the payoffs in this case with those that would obtain if funds were prohibited from open-ending. A first observation is that fund managers are always made worse off by going open-end: inspired managers get private benefits of $\theta Tz + (1 - \theta)Tz(1 + \alpha)$, and uninspired managers get z , both less than the values of $Tz(1 + \alpha)$ and Tz that they get, respectively, when they are restricted to being closed-end.

Perhaps more surprisingly, investors can also be made worse off by open-ending. With inspired managers pursuing a cautious investment strategy, their expected returns are given by $pq(\theta R^{SC} + (1 - \theta)R^{UC}) + (1 - pq)L$. Thus, we have

PROPOSITION 5. Suppose that conditions (1)–(4), as well as (7), all hold, so that in equilibrium funds are open-end, and inspired managers invest cautiously. A ban on open-ending would be Pareto-improving if, in addition, it is true that

(8)

$$pq(\theta R^{SC} + (1 - \theta)R^{UC}) + (1 - pq)L < pqR^{UC} + (1 - pq)K.$$

Observe that, from investors' perspective, the welfare comparison turns on a simple trade-off. On the one hand, open-ending allows for the early liquidation of uninspired managers; this yields an expected benefit of $(1 - pq)(L - K)$. On the other hand, open-ending forces inspired managers to invest too cautiously, thereby lowering their expected returns. Importantly, these are the only two factors. In equilibrium, open-ending does *not* confer any benefit in terms of improving the quality of the managerial pool. This is so even though the race toward an open-end form is

17. Note that Proposition 4 only sets out the conditions under which cautious investment is the *unique* equilibrium. For other parameter values, it is possible that there are *both* cautious-investment and aggressive-investment equilibria. This is because in a cautious-investment equilibrium, updating based on time-2 performance is harsher than in an aggressive-investment equilibrium: a performance of L or worse implies that the manager is uninspired for sure, as opposed to just being uninspired with probability $(1 - \pi)$.

motivated by an out-of-equilibrium desire on the part of good types to separate themselves from the pack.

If we set $p = 1$, so that there is no information asymmetry at time 0, then it becomes impossible to satisfy both (4) and (8) simultaneously. This means that any time the all-closed-end equilibrium is ruled out, investors are better off as a result of open-ending—in other words, the survival-of-the-fittest intuition of Fama and Jensen [1983] comes through. Thus, the model's welfare conclusions hinge on there being some ex ante heterogeneity in managerial quality, i.e., some bad types in the population. That such bad types exist is suggested by empirical studies which document the persistence of poor performance among the worst mutual funds.¹⁸ Moreover, as a matter of theory, it seems that some degree of heterogeneity would emerge endogenously in a fuller model, since potential managers with low ability will naturally be attracted by the closed-end form.¹⁹

EXAMPLE 1. Set $R^{UC} = 3.0$; $R^{SC} = 1.46$; $L = 0.95$; $K = 0.60$; $x = 0$; $\gamma = 0.1$; $\alpha = 0.05$; $z = 0.1$; $m = 3$; $T = 5$; and $p = q = 0.5$. It is easy to check that all six sufficient conditions, (1)–(4) and (7)–(8), are satisfied. Thus, all funds are open-end, and inspired managers put a fraction $\theta = 0.95$ of their portfolios in the SC asset. Investors earn an expected gross return of 1.10 for every dollar invested, inspired managers get private benefits of 0.501, and uninspired managers get private benefits of 0.10. By contrast, if there were a prohibition on open-ending, investors would earn an expected gross return of 1.20, inspired managers would get 0.525, and uninspired managers would get 0.50.

EXAMPLE 2. Keep everything else the same as in Example 1, but set $x = 0.90$. The first five conditions, (1)–(4) and (7), continue to be satisfied, so all funds are again open-end and inspired managers still invest cautiously. But now caution

18. See, e.g., Carhart [1997]. This literature suggests is that it is hard to find managers who consistently beat the market after fees, but that there are some consistent and quite glaring underperformers.

19. To be precise, consider a situation in which it is efficient for funds to be closed-end. If there is initially no heterogeneity among existing managers ($p = 1$), the all-closed-end equilibrium will be viable. But if there are any bad types in the general population, it would seem that they would be drawn into the fund sector by the prospect of enjoying undisciplined private benefits over a long horizon in the closed-end form. This will tend to push down p , generating the mixture of good and bad types that I assume.

only requires that a fraction $\theta = 0.50$ of their portfolios be allocated to the SC asset. As a result, investors earn an expected gross return of 1.27, which is better than in the all-closed-end scenario.

Overall, the two examples suggest that competition among funds can lead to an amount of open-ending that is excessive from investors' perspective, but not necessarily so.²⁰ In the first example, open-ending makes inspired managers so cautious that they almost completely shun the UC asset. The negative effect on their returns thus overwhelms the gains associated with being able to liquidate uninspired managers. In the second example—where the short-run downside risk of the UC asset is not as severe—open-ending does not induce quite so much caution, so the trade-off goes in the other direction. In the latter case, allowing funds to compete on the dimension of organizational form has the sort of welfare consequences that one might normally expect. That is, it makes the fund managers themselves worse off, but it benefits investors.

It should also be emphasized that in either case, open-ending dramatically reduces funds' willingness to invest in the UC asset. If we think of the interval from time 1 to time 3 as being five years, the numbers in the examples imply that the UC asset offers a *riskless* annual return of 25 percent if held to maturity, as compared with 8 percent for the SC asset. Yet in the first example, funds put only 5 percent of their portfolios in the UC asset; even in the second example, the portfolio weight is only 50 percent. Thus, even when open-ending can be said to be an efficient outcome, it greatly compromises the ability of arbitrageurs to take on certain types of trades.

One implication of this is that the connection between arbitrageurs' profits and market efficiency is tenuous. In Example 1 the gross return to investors of 1.10 translates into a 2 percent annual alpha for professionally managed money, assuming a

20. The welfare comparisons in Examples 1 and 2 are actually quite narrow in scope, as they do not consider the potential consequences that open-ending has for either market efficiency per se, or resource allocation. This is an artifact of the modeling strategy of taking the returns on the UC and SC assets to be exogenous. It seems likely that, in a fuller model with endogenous asset prices, the social costs of open-ending would loom even larger. In particular, open-ending would now tend to lead to more mispricing of UC-type assets, and to the extent that asset prices affect production decisions, to a less efficient allocation of resources, e.g., to too much investment in overpriced dot-com companies.

five-year horizon. This looks relatively small, in keeping with the empirical evidence on the performance of fund managers. Yet this small alpha coexists with an infinitely elastic supply of the UC asset, which can be thought of as having a price that deviates from fundamental value by a factor of three. Indeed, fund managers barely touch the UC asset, even though eventual convergence to fundamentals is assured, and there are no other frictions, such as trading costs or short-sales constraints. As noted in the introduction, this kind of story would seem to fit well with the unwillingness of hedge funds to bet heavily against the dot-com bubble of the late 1990s, by, e.g., taking an outright short position on the NASDAQ index—something that would certainly have been feasible to do at low cost from an execution/trading-frictions standpoint.

III.C. Performance-Based Incentives

One limitation of the model is that I have not considered the possibility that investors might put fund managers on an explicit incentive scheme.²¹ The only source of incentive alignment has come from the α parameter, which measures how much more in the way of private benefits managers get when they are able to maintain a position in the UC asset until time 3. And as can be seen from condition (7), it is crucial for the results of Propositions 4 and 5 that α not be too large. For if α is big enough, an inspired manager will continue to invest everything in the UC asset even when faced with meaningful liquidation risk in the open-end form. In this case, being open-end dominates being closed-end, and the model has little interesting to say.

Thus, once we allow for the possibility of incentive contracts, the results of Propositions 4 and 5 will only continue to apply if investors find it optimal to not make them too high-powered. In Stein [2004] I provide a formal analysis of this issue. I demonstrate that, under some conditions, overly high-powered incentives lead to value-destroying risk-shifting behavior on the part of uninspired managers. As a result, investors may endogenously choose to avoid such contracts, in which case the previous results continue to apply as stated.

21. That is, the model looks to be more an account of the mutual-fund sector, where there are no performance-based incentives, than of the hedge-fund sector.

V. CONCLUSIONS

Asset-management intermediaries such as mutual funds and hedge funds compete for investors' dollars, and one key dimension on which they compete is the choice of organizational form. In general, there is no reason to believe that this competition results in the form that is best-suited to the task of arbitrage. Rather, there is a tendency toward too much open-ending, which leaves funds unable to aggressively attack certain types of mispricing; i.e., those that do not promise to correct themselves quickly and smoothly.

This idea may help to shed light on the arbitrage activities of another class of players: nonfinancial firms. Baker and Wurgler [2000] document that aggregate equity issuance by nonfinancial firms has predictive power for the stock market as a whole. And Baker, Greenwood, and Wurgler [2003] show that such firms are also able to time the maturity of their debt issues so as to take advantage of changes in the shape of the Treasury yield curve. At first, such findings appear puzzling. After all, even if one grants that managers have some insight into the future of their own companies, and hence can time the market for their own stock, it seems harder to believe that they would have an advantage over professional arbitrageurs in timing the aggregate stock and bond markets.

However, there is another possible explanation for these phenomena. Even if managers of nonfinancial firms are less adept at recognizing aggregate-market mispricings than are professional money managers, they have an important institutional advantage—they conduct their arbitrage inside closed-end entities, with a very different incentive structure. For example, it was much less risky for the manager of an overpriced Internet firm to place a bet against the dot-com bubble in 1999 by undertaking an equity offering than it was for a hedge-fund manager to make the same sort of bet by, e.g., shorting the NASDAQ index. In the former case, if the market continued to rise, no visible harm would be done: the firm could just sit on the cash raised by the equity issue, and there would only be the subtle opportunity cost of not having timed the market even better. There would certainly be no worry about investors liquidating the firm as a result of the temporary failure of prices to converge to fundamentals.

APPENDIX

Proof of Proposition 1

To refine away the all-closed-end pooling equilibria with the Cho-Kreps [1987] criterion, we require the following conditions. First, a bad type must not want to deviate to open-end even if this would lead him to be judged a good type. Second, a good type must prefer to deviate to open-end if this leads him to be judged a good type. Implicit in the second condition is a third—namely that investors will prefer to give their money to a deviator if they think he is a good type, even though he is now open-end, which may affect his investment strategy. The proof strategy is then to take any perfect Bayesian equilibrium in which all funds are closed-end, and show that for that equilibrium, the above three conditions hold. This will imply that no all-closed-end equilibrium is robust.

A complication is that even after we specify investors' beliefs about the time-0 quality of a deviator (i.e., that he is a good type), we still need to pin down their expectations regarding how the deviator will trade upon learning at time 1 that he is inspired. That is, we have to look at all possible continuation equilibria that may follow after a deviation in the first stage. There are two subcases to consider. First, investors might conjecture that, upon learning that he is inspired, a manager will be cautious, and invest a fraction just over θ of his portfolio in the SC asset, where $\theta = (L - x)/(1 - x)$, thereby ensuring that the portfolio always has a time-2 value of greater than L . So if investors see a performance of L or less, they will be sure that the manager is uninspired, and will liquidate, since this is the dominant strategy when facing an uninspired manager.

Alternatively, investors might conjecture that, upon learning that he is inspired at time 1, a manager will be aggressive, and invest everything in the UC asset. In this case, a realization of L or worse at time 2 does not prove that the manager is uninspired, since even an inspired manager can get unlucky with the UC asset. Still, the Bayesian update conditional on a realization of L or worse is that there is now only a probability of $q\gamma/(q\gamma + (1 - q))$ that the manager is inspired.

If investors liquidate, they get x when the manager turns out to be inspired, and L when he turns out to be uninspired. If not, they get R^{UC} when the manager turns out to be inspired, and K

when he turns out to be uninspired. Thus, liquidation of a poor performer at time 2 continues to be optimal if

$$(A.1) \quad q\gamma R^{UC} + K(1 - q) < q\gamma x + L(1 - q).$$

Inequality (A.1), which can be rearranged to yield condition (1) in the text, thus ensures that no matter what investors believe about the investment policy of an inspired manager—that is, in any continuation equilibrium—they will liquidate any fund when its performance is less than or equal to L .

Assume that the condition holds. We can now entertain time-0 deviations from an all-closed-end equilibrium. Again, the deviation to be considered is one in which a fund announces that it is open-end, and continues to charge no fees. Supposing that investors interpret such a deviation as a signal that the manager is a good type, we are interested in establishing the incentives of both the good and the bad types to actually go through with the deviation. Consider first a bad type deviating to open-end. If he is judged to be good, the best that can happen is that he will get m dollars of money to manage. But he will only have this money under management until time 2, since it is certain that he will have an interim return of L and hence be liquidated. Thus, the best a bad type can get by deviating to open-end is mz . In contrast, we have seen that he gets Tz in the all-closed-end equilibrium. So a sufficient condition to deter a bad type from deviating is that $m < T$, which is condition (2) in the text.

Consider a good type deviating to open-end, and assume that he gets m dollars to manage. Once he has the money, the good type will learn at time 1 whether he is inspired. If he is, he can invest either cautiously or aggressively, and will pick whichever one yields higher expected private benefits. Thus, a lower bound on his private benefits can be obtained by assuming that he invests cautiously. If a good type invests cautiously when he is inspired, ex ante expected payoffs can be calculated as follows. With probability q , investors leave their money in until time 3 and get $(\theta R^{SC} + (1 - \theta)R^{UC})$, while the manager gets $m\theta Tz + m(1 - \theta)Tz(1 + \alpha)$. With probability $(1 - q)$, investors pull the plug at time 2 and get L , while the manager gets mz . Putting it together, the expected return to investors in this case is $q(\theta R^{SC} + (1 - \theta)R^{UC}) + (1 - q)L$, while the expected payoff to the good manager is $q(m\theta Tz + m(1 - \theta)Tz(1 + \alpha)) + (1 - q)mz$.

A sufficient condition for a good manager to prefer deviating is that his expected payoffs from doing so exceed those he gets in

the all-closed-end pooling case. This boils down to the requirement that

$$(A.2) \quad q(m\theta Tz + m(1 - \theta)Tz(1 + \alpha)) \\ + (1 - q)mz > qTz(1 + \alpha) + (1 - q)Tz.$$

It is easy to see that condition (3) in the text implies (A.2). Indeed, (3) is a simplified version of (A.2) which sets $\theta = 1$. In other words, as stated in the text, condition (3) is excessively conservative because it implicitly assumes that an inspired manager invests everything in the SC asset, instead of just a fraction θ of his portfolio.

So far, we have shown that the only reasonable beliefs in any all-closed-end equilibrium are that a deviator is a good type. Moreover, we have proved that a good type is willing to deviate, provided that investors give him enough money to manage. It remains to check that investors would actually prefer to give their money to a fund that deviates to open-end, assuming that such a fund is taken to be a good type. Recall that if they invest in the all-closed-end pool, investors get an expected return of $pqR^{UC} + (1 - pq)K$. Investors are obviously better off with a good deviator if, when he is inspired, the deviator invests aggressively—in this case, not only is the manager better, the investment policy is unchanged, and investors gain the right to liquidate.

Less clear-cut is the case where an inspired manager invests cautiously. As we have seen, investors' expected return in this case is $q(\theta R^{SC} + (1 - \theta)R^{UC}) + (1 - q)L$. So a sufficient condition for investors to give their money to a deviator is that

$$(A.3) \quad q(\theta R^{SC} + (1 - \theta)R^{UC}) + (1 - q)L > pqR^{UC} + (1 - pq)K.$$

Condition (4) in the text implies (A.3). As with condition (3), condition (4) as stated is excessively conservative, in that it sets $\theta = 1$.

It follows that under the conditions of Proposition 1, there are no equilibria with reasonable beliefs in which all funds remain closed-end.

Proof of Proposition 3

The all-open-end pooling equilibrium can be supported by the out-of-equilibrium belief that a deviator to closed-end is a bad type—with this belief, nobody will want to deviate from the equilibrium, because they would get no money to manage. More-

over, this belief survives our extension of the Cho-Kreps [1987] refinement. To see why, suppose to the contrary that investors believed that a deviator to closed-end was a good type. If, with this optimistic belief, a good type would ever choose to deviate, a bad type would always deviate as well. Thus it is not “unreasonable,” in the sense of Cho and Kreps, to think of a deviator as a bad type.

To be precise, note that a good type will deviate to closed-end only if such a deviation lands him m dollars of money to manage. But if this is true, then a bad type will also get m dollars by deviating, and will, in the closed-end form, be assured of private benefits of mTz . In contrast, by remaining in the open-end pool, a bad type only gets one dollar to manage, and is liquidated for sure at time 2. (See Lemma 1 in the text.) So by remaining in the open-end pool, a bad type gets private benefits of just z . Thus, any time the good type benefits by deviating, the bad type does as well.

Proof of Lemma 1

In the aggressive-investment scenario, a manager whose performance at time 2 is less than or equal to L has a posterior probability π of being inspired, where π is given by equation (5) in the text. If investors liquidate, they get x when the manager turns out to be inspired, and L when he turns out to be uninspired. If not, they get R^{UC} when the manager turns out to be inspired, and K when he turns out to be uninspired. Thus, liquidation of a poor performer is optimal if

$$(A.4) \quad \pi R^{UC} + (1 - \pi)K < \pi x + (1 - \pi)L.$$

Using (5), it is straightforward to check that condition (1) implies (A.4). This establishes the lemma.

HARVARD UNIVERSITY AND NATIONAL BUREAU OF ECONOMIC RESEARCH

REFERENCES

- Agarwal, Vikas, Naveen Daniel, and Narayan Naik, “Flows, Performance and Managerial Incentives in Hedge Funds,” working paper, London Business School, 2004.
- Baker, Malcolm, and Jeffrey Wurgler, “The Equity Share in New Issues and Aggregate Stock Returns,” *Journal of Finance*, LV (2000), 2219–2257.
- Baker, Malcolm, Robin Greenwood, and Jeffrey Wurgler, “The Maturity of Debt Issues and Predictable Variation in Bond Returns,” *Journal of Financial Economics*, LXX (2003), 261–291.

- Bebchuk, Lucian A., "Asymmetric Information and the Choice of Corporate Governance Arrangements," working paper, Harvard Law School, 2002.
- Bolton, Patrick, and David Scharfstein, "A Theory of Predation Based on Agency Problems in Financial Contracting," *American Economic Review*, LXXX (1990), 93–106.
- Bolton, Patrick, and David Scharfstein, "Optimal Debt Structure and the Number of Creditors," *Journal of Political Economy*, CIV (1996), 1–25.
- Brunnermeier, Markus K., and Stefan Nagel, "Hedge Funds and the Technology Bubble," *Journal of Finance*, LIX (2004), forthcoming.
- Calomiris, Charles W., and Charles M. Kahn, "The Role of Demandable Debt in Structuring Optimal Banking Arrangements," *American Economic Review*, LXXXI (1991), 497–513.
- Carhart, Mark M., "On Persistence in Mutual Fund Performance," *Journal of Finance*, LII (1997), 57–82.
- Cerulli Associates, Inc., "Market Update: Hedge Funds," *The Cerulli Report Series*, 1999.
- Chevalier, Judith, and Glenn Ellison, "Risk Taking by Mutual Funds as a Response to Incentives," *Journal of Political Economy*, CV (1997), 1167–1200.
- Cho, In-Koo, and David M. Kreps, "Signaling Games and Stable Equilibria," *Quarterly Journal of Economics*, CII (1987), 179–222.
- Chordia, Tarun, "The Structure of Mutual Fund Charges," *Journal of Financial Economics*, XLI (1996), 3–39.
- Diamond, Douglas W., "Debt Maturity Structure and Liquidity Risk," *Quarterly Journal of Economics*, CVI (1991), 709–737.
- Diamond, Douglas W., and Raghuram G. Rajan, "A Theory of Bank Capital," *Journal of Finance*, LV (2000), 2431–2465.
- Diamond, Douglas W., and Raghuram G. Rajan, "Liquidity Risk, Liquidity Creation, and Financial Fragility: A Theory of Banking," *Journal of Political Economy*, CIX (2001), 287–327.
- Fama, Eugene F., and Michael C. Jensen, "Agency Problems and Residual Claims," *Journal of Law and Economics*, XXVI (1983), 327–349.
- Flannery, Mark J., "Asymmetric Information and Risky Debt Maturity Choice," *Journal of Finance*, XLI (1986), 19–37.
- , "Debt Maturity and the Deadweight Cost of Leverage: Optimally Financing Banking Firms," *American Economic Review*, LXXXIV (1994), 320–331.
- Hart, Oliver D., and John Moore, "Default and Renegotiation: A Dynamic Model of Debt," *Quarterly Journal of Economics*, CXIII (1998), 1–41.
- Investment Company Institute, *Mutual Fund Fact Book*, 2003.
- Ippolito, Richard A., "Consumer Reaction to Measures of Poor Quality: Evidence from the Mutual Fund Industry," *Journal of Law and Economics*, XXXV (1992), 45–70.
- Lamont, Owen A., and Jeremy C. Stein, "Aggregate Short Interest and Market Valuations," *American Economic Review Papers and Proceedings*, XCIV (2004), 29–32.
- Lee, Charles M. C., Andrei Shleifer, and Richard H. Thaler, "Anomalies: Closed-End Mutual Funds," *Journal of Economic Perspectives*, IV (1990), 153–164.
- Nanda, Vikram, M. P. Narayanan, and Vincent A. Warther, "Liquidity, Investment Ability, and Mutual Fund Structure," *Journal of Financial Economics*, LVII (2000), 417–443.
- Ross, Stephen A., "The Princeton Lectures in Finance, Lecture 2: Efficient Markets," working paper, Massachusetts Institute of Technology, 2001.
- Shleifer, Andrei, and Robert W. Vishny, "The Limits of Arbitrage," *Journal of Finance*, LII (1997), 35–53.
- Sirri, Erik R., and Peter Tufano, "Costly Search and Mutual Fund Flows," *Journal of Finance*, LIII (1998), 1589–1622.
- Stein, Jeremy C., "Why Are Most Funds Open-End? Competition and the Limits of Arbitrage," NBER, Working Paper No. 10259, 2004.