

Managerial Ability, Compensation, and the Closed-End Fund Discount

JONATHAN B. BERK and RICHARD STANTON*

ABSTRACT

This paper shows that the existence of managerial ability, combined with the labor contract prevalent in the industry, implies that the closed-end fund discount should exhibit many of the primary features documented in the literature. We evaluate the model's ability to match the quantitative features of the data, and find that it does well, although there is some observed behavior that remains to be explained.

THE ROLE OF DELEGATED PORTFOLIO MANAGERS in financial markets has puzzled financial economists for some time. Their high level of remuneration suggests that they add considerable value. Yet, while their important role as providers of diversification services is widely recognized, economists have had difficulty finding evidence of any additional value that might justify their compensation levels. It is certainly not hard to find examples of funds that seem to add no value. For example, the Homestead Nasdaq 100 Index Tracking Stock Fund (ticker: HNASX) is an open-end fund that charges investors 1.5% per year to track the Nasdaq 100 index. As of September 30, 2004, 99.42% of the fund was invested in a single security, the NASDAQ 100 exchange-traded fund (QQQQ). The remaining 0.58% was in cash. Investors in the fund therefore pay 1.5% per annum to invest in a single security they could buy themselves for no annual charge.¹

This lack of evidence of ability raises serious questions about why investors entrust their savings to active managers at all, and nowhere is this puzzle starker than in the closed-end fund sector. Here, most funds trade at a discount

*Berk is at U.C. Berkeley and the National Bureau of Economic Research. Stanton is at U.C. Berkeley. This material is based upon work supported by the National Science Foundation under Grant No. 0339834. The authors gratefully acknowledge helpful comments and suggestions from Greg Brown, Martin Cherkas, Josh Coval, Bruno Gerard, Rick Green, Steve Grenadier, David Musto, Matthew Richardson, Steve Ross, Jacob Sagi, Mark Seasholes, Laura Starks, Josef Zechner, and seminar participants at Arizona, California State University (Fullerton), Columbia, Emory, Houston, INSEAD, London Business School, Massachusetts Institute of Technology, Michigan, Michigan State, Northwestern, Norwegian School of Economics, Norwegian School of Management, Simon Fraser, Texas A&M, UC Berkeley, UC Irvine, UNC, UT Austin, Wharton, the 2003 UBC Summer Finance Conference, the 2004 WFA meetings, and the 2004 University of Oregon conference on Delegated Portfolio Management. We are grateful for financial support from the Fisher Center for Real Estate and Urban Economics.

¹ See "Homestead fund is far afield for indexer" by Chuck Jaffe (Boston Herald, May 17, 2005).

to their net asset value (NAV), though they are issued at a price at or above the NAV. Why would investors buy a closed-end fund at its IPO, knowing that it is likely to fall to a discount, when they could instead buy an open-end fund that is guaranteed always to trade at par? The inability of economists to answer this question has led many to conclude that investor irrationality is the only possible explanation.

Recently, researchers have begun to find more compelling evidence of managerial ability. For example, Chen, Jegadeesh, and Wermers (2000) find a significant gap between the long-horizon returns of stocks that mutual funds buy and of stocks that they sell, and Litov et al. (2005) find that stocks that funds buy earn significantly higher returns at the next earnings announcement than stocks that they sell. This evidence suggests that managerial ability might play an important role in explaining investment company behavior. Berk and Green (2004) explore this idea in a recent paper, and show that the existence of managerial ability can explain an important puzzle in the open-end fund literature, namely, the relation between fund returns and capital flows in and out of funds. In this paper, we continue this line of research by exploring the implications of managerial ability in the closed-end fund sector. We show that the existence of managerial ability can also explain much of the observed behavior of closed-end funds.

In our model, the discount is driven by the tradeoff between managerial ability and fees. Managerial ability adds value to the fund, so, if there were no fees, competitive investors would be willing to pay a premium over net asset value to invest in the fund. Fees subtract value from the fund, so, if managers had no ability, investors would only be willing to invest if they could buy shares in the fund at a discount. In the presence of both fees and managerial ability, the fund may trade at either a premium or a discount depending on whether fees or ability dominate.

The discount changes over time as investors change their beliefs about the manager's ability, and as his fee changes (which depends on his compensation contract). Under a fixed-pay long-term compensation contract, whereby the manager's fee is a constant fraction of assets under management and the manager commits not to leave the fund, the fund may move to either a premium or a discount as investors' expectations of the manager's ability change over time. Under a short-term contract, whereby the manager's pay continuously adjusts so that his fee always equals the value he adds (so he has no incentive to leave), the fund always trades exactly at par. However, neither compensation contract is used in practice: The long-term contract is not feasible, because the manager cannot credibly commit to stay with the fund if he receives a more lucrative outside offer, and the short-term contract is unattractive, because it forces the manager to bear all of his human capital risk.

Under a compensation contract that is closer to those used in practice, we find that the discount exhibits dynamics similar to those actually observed. As with the long-term contract, we assume managers cannot be fired, and are paid

a fixed fraction of assets under management.² If performance is bad, investors infer that they have a bad manager, that is, a manager who charges more in fees than the value he creates. Because the manager is entrenched, the fund trades at a discount.³ If performance is good, investors infer that they have a good manager, that is, a manager who adds more in value than the fees he charges, and the fund trades at a premium. Unlike the long-term contract, however, the manager soon captures the benefit he provides by negotiating a pay increase (by threatening to quit the fund for more lucrative employment elsewhere). Thus, premia are short lived. Most funds sell at a discount, and this discount is a function of the age of the fund: Funds are issued at their NAV (par), they subsequently fall into discount, and this discount disappears close to the open-end date, when the capitalized value destroyed by the manager becomes small.

In our model investors are fully rational. Consequently, they understand that, by agreeing to the labor contract initially, they are providing the manager with employment insurance. They charge for this insurance by lowering the manager's wage, so that it is lower than the value he is initially expected to add. This is consistent with the recent finding by Wermers, Wu, and Zechner (2005) that new managers earn substantial excess NAV returns. In expectation, the additional value added by the manager over his fees is exactly offset by the expected growth in the discount, and so investors get a fair return. In other words, in our model a fully rational investor is willing to buy a fund for its NAV at the IPO, even though she expects the fund to fall into discount.

The results in this paper are closely related to those in Berk and Green (2004). Because the price of an open-end fund is forced to equal NAV at the end of each day, investors react to changes in their beliefs about managerial ability by moving capital in and out of the fund, thereby driving the returns of the fund to the competitive level. Berk and Green (2004) exploit this mechanism to explain the observed relation between open-end fund flows and performance. With closed-end funds, the assets under management remain fixed, so investors' updates of managerial ability cause *price* changes, which in turn drive returns to the competitive level. Taken together, the papers show that two different, widely recognized, empirical puzzles—the flow of funds—performance relation in open-end funds, and the behavior of closed-end fund discounts—both derive from the same economic fundamental, the existence of managerial ability.

² Past research has found that proxy contests are rare, and managers are very rarely fired. See, for example, Khorana (1996), Chevalier and Ellison (1999), Hu, Hall, and Harvey (2000), Khorana (2001), and Ding and Wermers (2004) for open-end funds, and Rowe and Davidson (2000) and Wermers, Wu, and Zechner (2005) for closed-end funds (as well as McIntosh, Rogers, and Sirmans (1994) for real estate investment trusts).

³ Brauer (1984, 1988) and Brickley and Schallheim (1985) demonstrate that very large gains are associated with managerial turnover, which supports the idea of entrenchment and suggests that large costs must be incurred in order to achieve the turnover. Brauer (1984) and Bradley et al. (2005) discuss the nature of these costs.

The paper is organized as follows. Section I reviews other explanations for the closed-end fund discount. Section II develops the formal model, and Section III calibrates the model, showing that it is qualitatively capable of matching the main empirical features of the closed-end fund discount. Section IV looks at the empirical literature on the discount in more detail, and shows that the model is quantitatively consistent with most of its empirical regularities, but that a number of puzzling aspects of the discount remain to be explained. Section V concludes the paper.

I. Prior Explanations for the Closed-End Fund Discount⁴

Spiegel (1999) shows that, even in a frictionless production economy, the price of a closed-end fund need not track its net asset value. In that paper, arbitrageurs have finite lives. This allows for the possibility of multiple rational expectations equilibria. In some of these, even portfolios paying zero dividends with certainty do not trade at a zero price. Thus, closed-end funds need not have the same value as their underlying assets, and, on average, trade at a discount. However, neither managerial ability nor management compensation, the focus of our paper, is explicitly modeled.

The idea that a closed-end fund should trade for a discount if the manager charges fees (but does not add value) was originally proposed by Boudreaux (1973). If managers charge fees, and provide nothing of value in return, then the value of the fund to investors should be lower than the fund's NAV. Gemmill and Thomas (2002) and Ross (2002a) show that if the fund pays out a fraction γ to investors each year and pays fractional management fees of δ each year, the discount is⁵

$$\frac{\delta}{\gamma + \delta}.$$

In particular, if the payout rate to investors is zero, the discount is 100% *regardless* of how small the fractional fee paid to managers each year is.⁶ Empirically, Malkiel (1977) does not find that fees significantly explain variation in the level of the discount. Kumar and Noronha (1992) find that differences in fees do explain a small proportion of the cross-sectional variation in discounts. The main drawback of this explanation is that it cannot alone explain why closed-end funds trade at a discount, while open-end funds with similar fees trade at NAV. Nor can it explain why closed-end funds are issued at or above their NAV.

⁴The discussion in this section benefited from the very clear surveys of the literature by Dimson and Minio-Kozerski (1999), Lee, Shleifer, and Thaler (1990), Anderson and Born (2002), and Dimson and Minio-Paluello (2002).

⁵This result first appears as the infinite horizon limit of a formula in Ingersoll (1976) (p. 89).

⁶A simple way to see this intuitively is to think of the fund manager as being awarded a fraction δ of the shares remaining in the fund each year. After t years, investors are left with $(1 - \delta)^t$ times the number of shares they started with, which goes to zero as $t \rightarrow \infty$.

We argue in this paper that it is not variation in fees, but rather variation in managerial ability that explains the cross-sectional variation in discounts.

If some managers add value, a fund will trade at a discount if investors believe it has a bad manager (so they do not make back their fees), and at a premium if investors believe the fund has a good manager. Lee et al. (1990) point out that for this to explain the prevalence of discounts, together with the premium at the IPO, investors must expect superior returns at the IPO, but then (predictably) expect poor performance later. Although Lee et al. (1990) cite this observation as evidence of investor irrationality, our model predicts exactly this behavior for the return on the fund's underlying assets. However, when this is combined with the time-series behavior of the discount, investors in the fund's shares always receive the fair rate of return, consistent with the results of Sias, Starks, and Tiniç (2001).

The three explanations given in Ross (2002b), Ferguson and Leistikow (2004), and Arora, Ju, and Ou-Yang (2003) are closely related to ours in that they attempt to simultaneously explain both the issue price and the subsequent discount. Like us, Ross derives a dynamic rational model of closed-end funds that explains the post-IPO discounts as a function of the difference between the value added by the manager and the fees charged. Initially, on the IPO, investors are asymmetrically informed about the quality of the manager, and because of the idiosyncrasies of the IPO process, this information is not revealed until after the issue. Hence, the fund can be issued at NAV. Although we derive many of the same implications, Ross's model differs from ours in a number of important respects. First, we do not appeal to an information asymmetry; rather, all of our participants are symmetrically informed. Our explanation for why the fund is issued at NAV does not rely on the idiosyncrasies of the IPO process itself. Second, Ross (2002b) does not examine the implication of a long-term labor contract on the fund's performance. Consequently, that paper cannot explain the *predictable* increase in the discount after the IPO, an important contribution of this paper. Third, Ross does not model the intertemporal rational expectations equilibrium, and thus does not make an inference about the dynamic behavior of the post-offering price relative to the pre-offering price. However, to be consistent with a rational expectations equilibrium, investors cannot be fooled, so their expectation of the post-offering price must be realized on average. In Ross's model this implies that to induce them to participate, on average the post-offering price must equal the pre-offering price. Ferguson and Leistikow (2004) develop a closely related model.

Arora et al. (2003) use a two-period model to explain the behavior of the closed-end fund discount. Although they use a different mechanism from ours (managers have inside information that they exploit), they show that under certain conditions, closed-end funds will be issued at a premium and, with certainty, will fall into discount. However, that paper cannot explain the wide variation in the discount across funds, nor can it explain the time-series behavior of the discount. In particular, there is no mechanism in the model that can explain why a fund that is trading at a discount could move to a premium.

Another class of proposed explanations for the closed-end fund discount is based on market frictions such as illiquidity and taxes. If a fund owns a lot of restricted stock or other illiquid assets that do not trade freely, its NAV may not accurately reflect its true value, in which case the fact that it does not trade at its NAV is not particularly surprising. Malkiel (1977) and Lee, Shliefer, and Thaler (1991) find that holdings of restricted stock do have some explanatory power for discounts, but these holdings are small or zero for most funds, so cannot fully explain the behavior of the discount. Seltzer (1989) suggests that funds holding illiquid assets are likely to be overvalued, but this is inconsistent with the fact that the price of a fund rises when it is opened.

Recently, Cherkes, Sagi, and Stanton (2005) reexamine the importance of liquidity in closed-end funds. Rather than focusing on possible mispricing, they develop a model in which the discount is determined by investors' tradeoff between fees paid and liquidity benefits that accrue through holding the closed-end fund, rather than holding the underlying assets directly. Like ours, their paper is successful at explaining many of the observed features of the discount, suggesting that liquidity is an important additional factor in closed-end funds.

Full taxes on a fund's realized capital gains are paid by current shareholders even if most of the gains occurred before they bought their shares. This would imply that funds with large accumulated gains should trade at a discount to NAV. However, Malkiel (1977) finds that even a fund with (a very high) 25% of its assets in unrealized appreciation would see an average discount of only 5%, and moreover, the fact that prices rise to NAV upon fund liquidation suggests that this factor cannot be the main factor driving discounts.

Brickley, Manaster, and Schallheim (1991) and Kim (1994) suggest an alternative tax timing explanation based on the idea that holding shares indirectly via a closed-end fund precludes an investor from directly trading the underlying shares necessary to follow the optimal tax timing strategy.⁷ The empirical evidence is mixed. For example, Kim (1994) documents a large increase in the number of closed-end funds after 1986, when changes in the tax law reduced the tax disadvantage of holding closed-end funds, but DeLong and Shliefer (1992) document that the discount *increased* between 1985 and 1990. In addition, the tax timing option cannot explain funds trading at a premium, and should apply to both open- and closed-end funds.

II. The Model

This section develops a formal model of closed-end fund management and investment. We assume that the fund's asset holdings are observable periodically, and that any investor can costlessly mimic this investment strategy by adjusting his portfolio as soon as the information about holdings is released. The instantaneous return on the fund's current portfolio is given by

⁷ See, for example, Constantinides (1983, 1984).

$$d\hat{R}_t = r dt + \nu dW, \tag{1}$$

where the assets have expected (risk-adjusted) return r and volatility ν . The return generated by the manager for the fund is given by

$$dR_t = d\hat{R}_t + \alpha dt + \sigma dZ. \tag{2}$$

In general, dR_t does not equal $d\hat{R}_t$, because the manager has the potential to add value by trading. These potentially value-enhancing trades add additional expected return α and volatility σ , where dZ and dW are assumed to be uncorrelated. If the manager’s skill level, α , is positive (negative), the manager earns an average return greater (less) than the expected return on the portfolio at the beginning of the period. The manager charges a proportional fee, c_t , per period, and the fund pays a proportional dividend at rate d , so the fund’s net asset value, N_t , evolves as follows:

$$\begin{aligned} dN_t/N_t &= dR_t - (c_t + d) dt \\ &= (r + \alpha - c_t - d) dt + \sigma dZ + \nu dW. \end{aligned} \tag{3}$$

Define the (observable) excess NAV return

$$\begin{aligned} dX_t &\equiv dN_t/N_t + d dt - d\hat{R}_t \\ &= (\alpha - c_t) dt + \sigma dZ. \end{aligned}$$

Consider a manager who begins managing the fund in period 0, and assume that neither the manager nor the investors know the value of α , the manager’s skill level. They both have the same prior on α at time 0—normal with mean ϕ_0 and precision $\gamma = 1/\eta^2$. After observing the returns the manager generates, investors and managers alike update their priors. Let ϕ_t denote the expectation of α conditional on seeing all returns up to and including r_t , that is,

$$\phi_t = E_t[\alpha]. \tag{4}$$

The evolution of the priors is governed as follows:

PROPOSITION 1: *The posterior distribution of α is normally distributed, with variance $\frac{1}{\gamma+t\omega}$ (where $\omega \equiv 1/\sigma^2$), and mean ϕ_t , which evolves jointly with X_t according to*

$$dX_t = (\phi_t - c_t) dt + \sigma dZ, \tag{5}$$

$$\begin{aligned} d\phi_t &= \frac{\omega}{\gamma + \omega t} [dX_t - (\phi_t - c_t) dt] \\ &= \frac{\sqrt{\omega}}{\gamma + \omega t} dZ. \end{aligned} \tag{6}$$

Proof: Follows directly from theorem 12.1 in Liptser and Shiriyayev (1978). Q.E.D.

Note that in general $c_t \equiv c(\phi_t, t)$ will be a function of both ϕ_t and t , the exact functional form depending on the details of the manager's labor contract. In addition, note that the future distribution of ϕ , and hence any function of ϕ , such as the discount, determined by the dynamics in equation (6), is the same under both the true and risk-neutral distributions.

A. The Discount

Assume the manager gets a fixed amount of capital to invest at time 0, and the assets in the fund are distributed to the shareholders (at net asset value) at the open-end date, T .⁸ Competition among investors implies that, in equilibrium, they cannot earn an expected return greater than r by investing in the fund. Because they are rational, they will not accept an expected return lower than r , so their expected return from investing in the fund must always exactly equal r . However, the expected return the manager generates, after fees, will in general not equal r , so the price of the fund adjusts, implying that it will not equal NAV. Let $P(N, \phi, t, T; d)$ be the price at time t of a fund whose NAV equals N , and define $D(\phi, t, T; d) \equiv P(N, \phi, t, T; d)/N$, so the fund's discount is $1 - D(\phi, t, T; d)$.

PROPOSITION 2: $D(\phi, t, T; d)$ is not a function of r ; rather, it satisfies the partial differential equation

$$D_t + (\phi - c(\phi, t) - d)D + \frac{1}{\gamma + t\omega} D_\phi + \frac{1}{2} \frac{\omega}{(\gamma + t\omega)^2} D_{\phi\phi} + d = 0. \quad (7)$$

Proof: By Ito's Lemma, we have

$$\begin{aligned} dP = & \left[P_t + P_N N(r + \phi - c(\phi, t) - d) + 1/2 P_{NN} N^2(\sigma^2 + \nu^2) + 1/2 P_{\phi\phi} \frac{\omega}{(\gamma + t\omega)^2} \right. \\ & \left. + P_{N\phi} \frac{N\omega\sigma^2}{\gamma + t\omega} \right] dt + \left[P_N N\sigma + P_\phi \frac{\omega\sigma}{\gamma + t\omega} \right] dZ + P_N N\nu dW. \end{aligned} \quad (8)$$

The expected return (including dividends) from investing in the fund must equal r , so

$$E_t(dP) = (rP - dN) dt. \quad (9)$$

Write

$$P = ND(\phi, t, T; d). \quad (10)$$

⁸ We impose a finite horizon to see the behavior of the discount both after the IPO and as the fund approaches a (known) open-end date. In practice, the open-end date is, of course, stochastic. If we assume that a fund can open up in any instant with probability $s dt$, and that the event of open-ending is uncorrelated with returns, everything to follow continues to hold as long as we set the dividend rate in all formulae to $d + s$ instead of d .

Equation (7) follows from equating the coefficients of dt in equations (8) and (9), substituting for derivatives of P from equation (10), and dividing through by N . Independence of r follows by inspection. Q.E.D

At the open-end date, T , the assets are distributed to investors, so the absence of arbitrage guarantees that D must satisfy the following boundary condition:

$$D(\phi, T, T; d) = 1 \quad \forall \phi. \tag{11}$$

The additional boundary conditions, and the level of compensation, $c(\phi, t)$, are determined by the manager’s labor contract. We shall consider three different labor contracts, each of which results in very different implications for the discount.

Long-Term Contract. We begin by assuming that the manager signs a *long-term* compensation contract that promises to pay him a constant fraction of the assets each period until time T , that is,

$$c(\phi, t) \equiv c \quad \forall \phi_t, t.$$

This contract is binding on both the manager and the fund (so the manager can neither quit nor be fired). Under this contract, D satisfies the following two additional boundary conditions:

$$\lim_{\phi \rightarrow \infty} D(\phi, t, T; d) = \infty, \tag{12}$$

$$\lim_{\phi \rightarrow -\infty} D(\phi, t, T; d) = 0. \tag{13}$$

It is straightforward to verify that the solution to equation (7), subject to boundary conditions (11), (12), and (13), is

$$D(\phi, t, T; d) = d \int_t^T \exp \left\{ \left(\phi - c - d + \frac{s - t}{2(\gamma + t\omega)} \right) (s - t) \right\} ds + \exp \left\{ \left(\phi - c - d + \frac{T - t}{2(\gamma + t\omega)} \right) (T - t) \right\}. \tag{14}$$

When there are no dividends, so $d = 0$, the solution is

$$D(\phi, t, T; 0) \equiv D_0(\phi, t, T) = \exp \left\{ \left(\phi - c + \frac{T - t}{2(\gamma + t\omega)} \right) (T - t) \right\}. \tag{15}$$

Note that equation (14) can be written in terms of D_0 :

$$D(\phi, t, T; d) = d \int_t^T e^{-d(s-t)} D_0(\phi, t, s) ds + e^{-d(T-t)} D_0(\phi, t, T). \tag{16}$$

This expression has an intuitive interpretation: We can think of the fund in the presence of dividends as a portfolio of nondividend-paying funds, one

maturing next period (with assets equal to the present value of next period's dividend), plus another maturing the following period (with assets equal to the present value of the following period's dividend), etc. Therefore, $D(\phi, t, T; d)$ is a weighted average of D calculated without dividends, but with varying maturities, leading to equation (16).⁹

It is clear from (14) and (16) that in general, $D(\phi, t, T; d) \neq 1$. Even if we pick c such that D is initially equal to 1, if the manager turns out to be better than expected the fund will trade at a premium, and if he turns out to be worse than expected it will trade at a discount.

Short-Term Contract. Ex ante, the contract described above is attractive to a risk-averse manager, since he does not bear any of the risk that he might later turn out to be of low ability. However, it is not a feasible contract in practice. The manager cannot be prevented from leaving, so if he turns out to be good, he will be able to negotiate a pay raise by threatening to quit for more lucrative employment elsewhere, regardless of what was agreed to ex ante.

An alternative contract that avoids this problem is one in which the manager's pay continuously adjusts so that it always exactly equals the value he adds (or subtracts) in every period. Under this *short-term* labor contract, $c = \phi_t$, so the manager is always paid what he is worth and has no incentive to quit, but he bears all his human capital risk. In this case, equation (7) becomes

$$D_t - dD + \frac{1}{\gamma + t\omega} D_\phi + \frac{1}{2} \frac{\omega}{(\gamma + t\omega)^2} D_{\phi\phi} + d = 0, \tag{17}$$

with the solution $D(\phi, t, T; d) = 1$. The fund always trades for its NAV.

Insurance Contract. In practice, actual closed-end fund managers do not work under either of the two contracts described above. Managers very rarely work under short-term contracts that force them to take on all their human capital risk, and the long-term contract is not feasible. Under the constraint that managers cannot be forced to work, Harris and Holmström (1982) show that the optimal managerial contract specifies that the manager works for an initial

⁹ The integral in equation (14) can be evaluated in closed-form, yielding

$$D(\phi, t, T; d) = 1/2 d \sqrt{2\pi(\gamma + t\omega)} e^{-1/2(\phi - c - d)^2(\gamma + t\omega)} \\ \times \left(\operatorname{erfi} \left(\frac{(T - t) + (\phi - c - d)(\gamma + t\omega)}{\sqrt{2(\gamma + t\omega)}} \right) - \operatorname{erfi} \left(\frac{(\phi - c - d)(\gamma + t\omega)}{\sqrt{2(\gamma + t\omega)}} \right) \right) \\ + e^{(\phi - c - d + \frac{T-t}{2(\gamma + t\omega)})(T-t)},$$

where erfi is the imaginary error function, which is defined by

$$\operatorname{erfi}(x) \equiv -i \operatorname{erf}(ix).$$

This expression is not very useful numerically, as for large arguments, we end up subtracting one huge value from another. In practice, it is better just to integrate equation (14) numerically.

wage that is never cut, but is increased any time the manager's outside options are better than his current contract by just enough to induce the manager to stay on. That is, employers insure their employees' human capital by contracting never to lower their wages (or fire them), but they continue to give them pay increases. Of course, frictions in actual labor markets prevent employers and employees from implementing this contract perfectly. For example, switching jobs usually entails costs, so employers do not have to raise employees' salaries until their outside options exceed these costs. We shall therefore consider a modified version of this contract that is a more realistic approximation of actual labor contracts in practice.

Assume that managers are paid a fixed fraction c of assets under management. When the manager's perceived ability, ϕ_t , rises above some level $\bar{\phi}$, he receives a more lucrative outside offer of employment and so is able to demand, and receive, a pay raise. The actual form this pay raise takes is not important. For simplicity, we use the mechanism most prevalent in (open-end) mutual funds, namely, we assume that the manager can raise additional (unmodeled) capital to manage, either within the existing fund or by starting an additional fund. In either case, he is spread thinner (his skill level is fixed but it must now be employed over a larger scope), which negatively impacts his ability to generate excess returns. As in Berk and Green (2004), we assume that the manager will raise capital up to the point that ϕ_t is driven to the level (denoted ϕ_t^*) at which the fund again trades at par.¹⁰ This adjustment to ϕ_t does not affect the confidence investors have in their estimates of the manager's ability, so (6) continues to describe the future evolution of ϕ_t . This assumption translates into the upper boundary condition

$$D(\bar{\phi}, t, T; d) = 1. \quad (18)$$

Although we assume that the manager cannot be fired (capital cannot leave the fund except through dividend payments), because any manager always has an option to index, we assume managers do not knowingly destroy value. Consequently, if the expectation of the manager's ability drops below zero (the manager expects to add no value), we assume he exercises this option and indexes from then on.¹¹ To derive the boundary condition that imposes this assumption, we need to calculate the discount at the point the manager decides to index. In this case D will still solve equation (7), but with ϕ and all derivatives with respect to ϕ set to zero (since there is no further updating):

$$D_t - (c + d)D + d = 0. \quad (19)$$

¹⁰ This is consistent with the results of Khorana, Wahal, and Zenner (2002), who study secondary rights offerings during the period 1988 to 1998. They find that rights offerings are announced, on average, when funds are trading at a premium, and that the premium disappears during the course of the offering.

¹¹ For simplicity, we ignore any real options that might imply waiting beyond this point.

The solution to this equation, with associated boundary condition

$$D(T, T; d) = 1 \quad \forall \phi, \quad (20)$$

is

$$D_{\text{index}}(t, T; d) = \frac{d + c e^{-(c+d)(T-t)}}{c + d}. \quad (21)$$

Thus, the lower boundary condition at $\phi = 0$ is

$$D(0, t, T; d) = \frac{d + c e^{-(c+d)(T-t)}}{c + d}. \quad (22)$$

Note that an equally plausible lower boundary condition is that investors incur the costs to fire the manager whenever ϕ_t drops to zero (or some other predetermined level). Then the discount will equal the cost of firing.

Finally, we set c , the manager's fee, so that the fund initially trades for its NAV:

$$D_0(\phi_0) = 1. \quad (23)$$

Henceforth we refer to this contract as the *insurance* contract. To determine the behavior of the discount under this contract, we solve (7) numerically, subject to the boundary conditions (11), (18), and (22), using an implicit finite-difference algorithm.

III. Empirical Implications of the Model

This section calibrates the model to the data using results from prior research. Our objective here is two-fold. First, we would like to characterize the effect of the labor contract on the behavior of the discount. This requires solving the partial differential equation numerically, and so we would like to use values that make sense. Second, we would like to evaluate whether the behavior of the discount in our model can reproduce the main characteristics of the discount observed in the data.

A. Calibration

We set the dividend payout rate to 2%, a number within the range of median dividend yields reported by Bradley et al. (2005) in their sample. Lee et al. (1990) report that management fees typically range between 0.5% and 2% per year. We use 1%. Next, in order to examine the model's implications for the discount both after the fund's IPO and as the fund approaches its open-end date, we impose a fixed terminal horizon of $T = 50$ years. This value ensures that, on average, only 2% of funds open-end per year, consistent with the low open-ending rates observed in practice.¹²

¹² Brauer (1984) documents a total of only 14 closed-end funds open-ending between 1960 and 1981. The frequency of open-ending has increased recently, but even so averages well under 5% per year in the period 1988 to 2003 (see Bradley et al. (2005)).

Table I
Parameter Values

Variable	Symbol	Value
Percentage fee	c	1%
Dividend	d	2%
Years to open-end date	T	50
Mean of prior	ϕ_0	3.4%
Prior precision	γ	550
Return precision	ω	40
Pay raise level of ability	$\bar{\phi}$	3.5%

Next, we turn to ω , the precision of the difference between the return of the fund and the return of the underlying assets held at the beginning of the period. Using U.S. data (1965 to 1985) Pontiff (1997) reports a mean monthly variance across funds of the difference between these returns equal to 37.33. Adams (2000) finds a value of 11.50 for U.K. funds (1982 to 1996). As Adams (2000) points out, Pontiff's data cover a particularly volatile period. More recently the United States has been much closer to the U.K. experience. Furthermore, the distribution of funds' volatilities is highly skewed—Pontiff reports a median of only 19.62. In light of this we set $\omega = 40$. Expressed on a monthly basis for percentage returns, this translates into a variance of about 20, somewhat less than Pontiff's mean, but larger than his median.

The remaining three parameters, the prior mean, ϕ_0 , the precision, γ , and the level at which the manager gets a pay raise under the insurance contract, $\bar{\phi}$, are not directly observable. We infer their values by experimenting with parameters that give realistic values for the average discount and maximum premium under the insurance contract. Table I summarizes the final parameters.

Note that these values of ϕ_0 and γ imply, at least initially, that 71% of the managers have sufficient skill to make back what they charge in fees. Because the average manager initially has an α of 3.4% but only charges a 1% fee, his initial NAV excess return of 2.4% is substantially higher than the fees he charges, but is close to the 1.94% NAV excess return for new managers reported in Wermers et al. (2005).

B. Discount vs. Ability

We begin by solving for the discount as a function of ability. We have closed-form expressions for both the long-term and short-term contracts. For the insurance contract, we solve (7) numerically. Figure 1 plots D against ability, ϕ , for new funds and for funds aged 10, 20, 30, 40, and 49 years. The solid lines describe the behavior of D under the insurance contract, and the dotted lines show the behavior under the long-term contract.¹³ Figure 2 plots discount against

¹³ Under the short-term contract, D is always equal to 1 and is therefore not shown.




Figure 1. Discount as a function of ability. The plots show the prices (expressed as % NAV) of new funds and funds of ages 10, 20, 30, 40, and 49 years, as a function of perceived management ability (ϕ). All funds have an initial open date of 50 years. The solid lines are funds with the insurance contract, and the dotted lines are funds with the long-term contract.

ability under the insurance contract and the long-term contract for funds of all ages.

The plots make clear that the behavior of the discount depends critically on the manager's compensation contract. Under the short-term contract, $D = 1$ and so there is no relation between ability and the discount. At the other extreme, under the long-term contract, D is monotonically increasing in ability: Funds with good managers trade at large premia, and funds with bad managers trade for discounts. Under the insurance contract, for low values of ϕ , the relation between ability and the discount is monotonic—a positive change in ability (a high return) translates into higher expected returns, which lead in turn to smaller discounts (D increases). This relation is documented empirically by Chay and Trzcinka (1999). However, when the fund is trading at a premium, in contrast to the long-term contract, at high enough levels of ϕ , the discount and ability are *inversely* related, so an increase in ability translates into a *lower* value of D (or, equivalently, a decrease in the premium). In this case, greater ability does not translate into higher long-term returns because there is a higher likelihood that the manager will receive a pay raise, either now or in the near future (when ϕ hits $\bar{\phi}$), causing the fund to trade for par.

The existence of a region in which $D(\cdot)$ is decreasing is interesting because it may help explain an empirical phenomenon that past researchers have found puzzling. Lee et al. (1991) and Ross (2002b) find that the empirical relation between past performance and the discount is negative—when funds outperform the market, the discount subsequently increases. As Figure 1 makes clear, good current performance need not translate into future outperformance if the result is that the manager gets a pay raise.

For the same level of compensation, when investors provide human capital insurance they demand much higher initial qualifications. Figure 3 shows ϕ_t^* , the value of ϕ for which the fund trades at par (or the initial qualification that

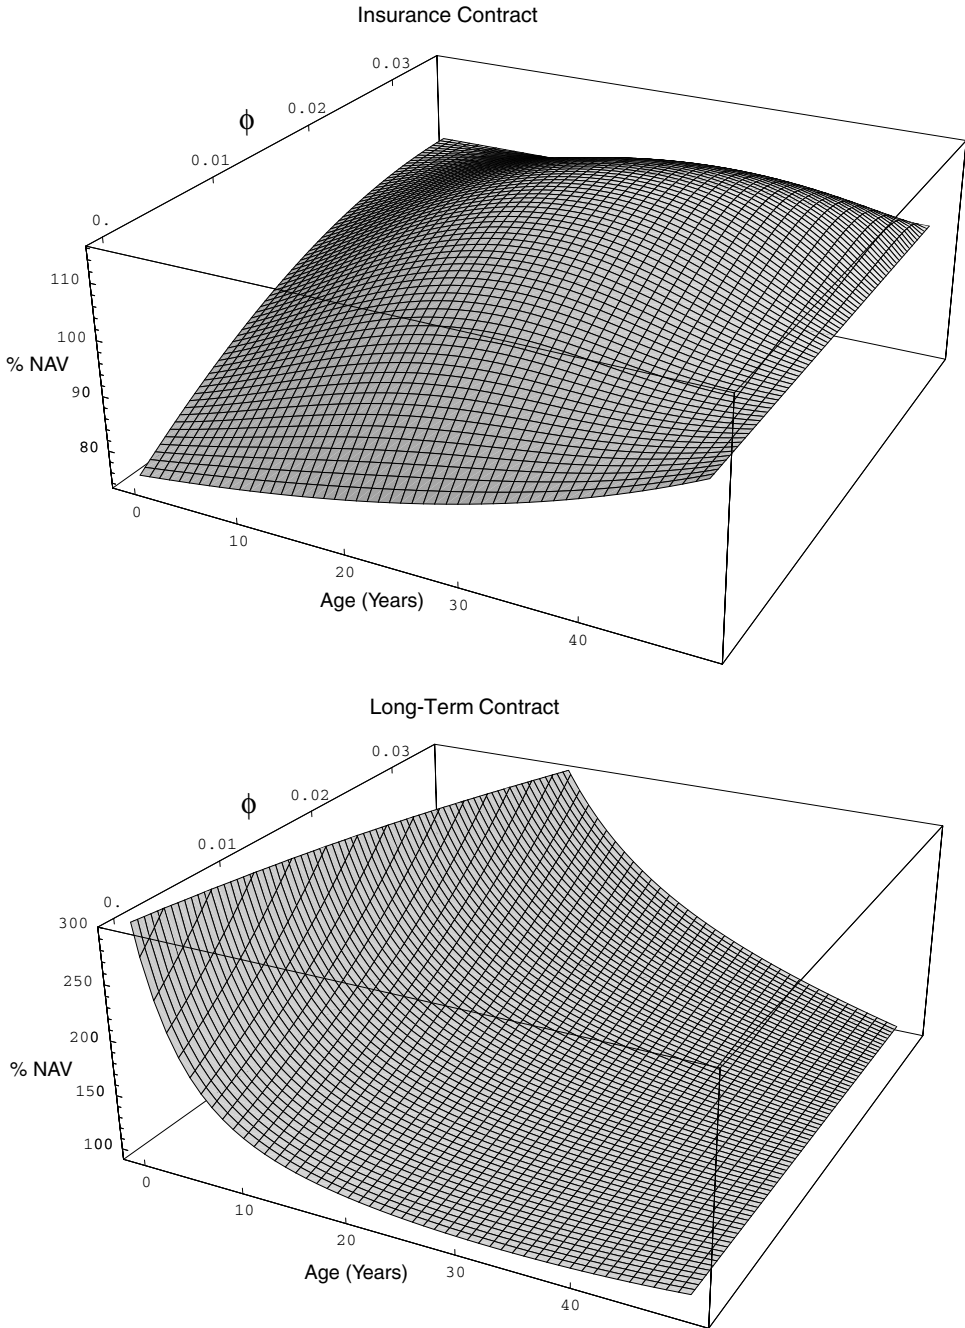


Figure 2. Discount as a function of ability. The two plots are the price of the fund under the insurance contract and the long-term contract, respectively, expressed as % NAV, for funds with an open date of 50 years as a function of perceived management ability (ϕ) and the age of the fund. Note that the vertical scales on the two plots are different.




Figure 3. Ability of the manager when the fund trades for NAV. The solid line shows ϕ_t^* , the value of ϕ for which the fund trades at par (NAV) under the insurance contract, as a function of the age of the fund. The dashed line is the same plot, but for a fund under the long-term contract.

allows the fund to be issued for NAV) under the long-term contract and the insurance contract (the fund always trades for par under the short-term contract). Under the insurance contract, investors initially require that managers have ability $\phi_0 = 3.4\%$, while under the long-term contract, they only require that $\phi_0 = -2.9\%$,¹⁴ an initial difference in ability of over 6%. This difference shrinks as the number of years to the open-end date decreases both because the ability of the manager becomes better known, and because the remaining length of time that investors must insure the manager decreases.

From Figure 2 it is clear that neither the long-term nor the short-term contract produces discount behavior that matches the behavior observed. Under the short-term contract the fund always trades at par, while under the long-term contract, a substantial number of funds, especially in their early lives, will have very large premia.

C. Expected Evolution of the Discount

One of the most puzzling open questions in the closed-end fund literature is why anybody would purchase a closed-end fund on its initial public offering, expecting the fund to fall into discount (see, e.g., Lee et al. (1990)). In our

¹⁴ Investors are willing to hire such a low-quality manager because of the convex relation between value and ability. Since the manager cannot quit, the disproportionate gain if the manager turns out to be good makes up for the fact that investors expect, on average, that the manager will destroy value.

model, the existence of managerial ability delivers exactly this result. Initially, managers charge less than the value they add, $\phi_0 > c$. Although they deliver this difference to investors, the average fund falls into discount because good managers extract pay raises while bad ones maintain their old compensation and become entrenched. The growth in the discount exactly offsets the value delivered to investors, leaving investors with the fair market return.

To derive the behavior of the expected discount, define

$$E(\phi, t, s, T; d) \equiv E_t[D(\phi_s, s, T; d) \mid \phi_t = \phi],$$

the expected price of the fund (expressed as % NAV) at a future time s , where $t \leq s \leq T$. By iterated expectations,

$$E(\phi, t, s, T; d) = E_t[E(\phi, \tau, s, T; d)],$$

for $t \leq \tau \leq s$, and hence

$$E_t(dE) = 0. \tag{24}$$

By Ito's Lemma, this can be written (dropping the arguments for clarity) as

$$E_t + 1/2E_{\phi\phi} \frac{\omega}{(\gamma + t\omega)^2} = 0. \tag{25}$$

The boundary condition at time s is

$$E(\phi, s, s, T; d) = D(\phi, s, T; d).$$

Under the insurance contract, the other boundary conditions are

$$E(\bar{\phi}, t, s, T; d) = E(\phi_t^*, t, s, T; d), \text{ and}$$

$$E(0, t, s, T; d) = \frac{d + c e^{-(c+d)(T-s)}}{c + d}.$$

The solid line in Figure 4 is the numerical solution of this partial differential equation as a function of the age of the fund. The fund is issued at par, but immediately after issuance it is expected to trade at a discount. The discount initially increases as more managers become entrenched and ultimately turn to indexing. Eventually, though, as the fund approaches its open-end date, the discount decreases, and it disappears completely on the open-end date. This is because, as the fund approaches its open-end date, the total capitalized value of what these poor managers expect to dissipate declines.

D. Distribution of the Discount

The wide cross-sectional variation in discounts that has puzzled researchers is an implication of our model. To quantify the distribution of the discount at some future date we first calculate the distribution of the manager's ability, and




Figure 4. Discount as a function of fund age. The solid line is the expected price of the fund, expressed as % NAV, as a function of fund age. The dashed lines show different fractiles of the distribution of the price of the fund expressed as % NAV, as a function of the age of the fund. The fractiles shown are marked on the graph.

then invert the relation between the discount and ability given by the solution to equation (7).

Let $f(\phi, t)$ be the probability density function of the manager’s ability, ϕ , at some future date, t , conditional on starting at date 0 with initial ability ϕ_0 . We calculate f by solving the Kolmogorov forward equation numerically.¹⁵ Note that $D(\phi, t, T; d)$ is not monotonic in ϕ , so it does not have a unique inverse. However, as Figure 2 illustrates, we can split the range of ϕ into two subsets, on each of which $D(\phi, t, T; d)$ is monotonic in ϕ . Let $\hat{\phi}_t$ be the value of ϕ at which $D(\phi, t, T; d)$ reaches a maximum, let $\phi_t(D)$ denote the inverse of $D(\phi, t, T; d)$ over the region $\phi \in (-\infty, \hat{\phi}_t]$, and let $\bar{\phi}_t(D)$ denote the inverse of $D_t(\phi_t)$ over the region $\phi \in (\hat{\phi}_t, \bar{\phi})$.

The following proposition derives the distribution of the discount as a function of management tenure. The proof is straightforward and left to the reader.

PROPOSITION 3: *Let $g_t(D)$ be the probability density function of D , the price of the fund expressed as a fraction of net asset value, for a fund at time t . Then*

$$g_t(D) = \begin{cases} f(\phi_t(D), t) \frac{d}{dD} \phi_t(D) & -\infty < D < 1 \\ f(\phi_t(D), t) \frac{d}{dD} \phi_t(D) - f(\hat{\phi}_t(D), t) \frac{d}{dD} \hat{\phi}_t(D) & 1 \leq D < D(\hat{\phi}_t, t, T; d). \end{cases}$$

¹⁵ See Øksendal (2002).

Figure 4 plots the 95%, 90%, 75%, 50% (the median), 25%, 10%, and 5% fractiles of g using the parameters in Table I. For example, 95% of funds are expected to have a value of D below the top line in the figure. It is clear from the figure that the model generates a wide dispersion in discounts. After 5 years, the top 25% of funds are expected to trade at premia in excess of approximately 1% even while the bottom 25% of funds are expected to trade at discounts of approximately 25%. After about 15 years, 25% of funds are still expected to trade at a premium, even while more than 50% of funds are expected to trade at discounts of more than 20%. The model has a lower absorbing barrier at zero (any time ϕ drops to zero, the manager indexes). Eventually, therefore, every fund will end up indexing, which explains why the lines converge to a lower bound, the value of D when the manager is indexing. Even so, it takes almost 10 years for the majority of funds to index, and almost 20 years for 75% of funds to index.

The overall level of the discount is within the range observed in the data. Lee et al. (1991) analyze U.S. funds between 1965 and 1985, and report that discounts towards the end of this period tend to be between 10% and 20%. From 1980 to 1998, the average discount on U.S. stock funds varied from about zero to 20%,¹⁶ and discounts at the end of the 1990s were around 10% in the United Kingdom and 5% in the United States. Anderson and Born (2002) report that in February 2001 the average discount for all equity funds (worldwide) was 10.9%.

E. NAV Return

Lee et al. (1990) point out that, in order for a model with differential management ability to generate the observed time-series behavior of average discounts, investors must systematically expect management performance to decline over time. Our model predicts exactly this behavior.

The expected excess NAV return, $E_0[\phi_t] - c$, is straightforward to calculate given $f(\phi, t)$. Figure 5 plots this expectation as a function of fund age, t . It also plots the 95%, 90%, 75%, 50% (the median), 25%, 10%, and 5% fractiles of distribution of the NAV excess return as a function of fund age. The manager is initially expected to outperform the benchmark (the return on the assets held at the beginning of the period), but as the tenure of the manager increases, he is eventually expected to underperform. Because he can do no worse than simply indexing, his 1% fee represents a lower bound on the value he adds. After 19 years, 75% of funds are at the lower bound, that is, the manager is expected to destroy 1% per year. Even so, 10% of funds are still expected to generate at least 75 basis points per year after 19 years.

¹⁶ See Dimson and Minio-Kozerski (1999).




Figure 5. NAV excess return as a function of fund age. The solid line is the instantaneous expected NAV excess return, $E_0[\phi_t] - c$, as a function of fund age, t . The dashed lines are the different fractiles of the distribution of NAV excess return, $\phi_t - c$, as a function of fund age, t . The fractiles shown are marked on the graph.

IV. Comparison with Existing Empirical Results

Our objective in this paper is to show that by explicitly modeling managerial ability, it is possible to construct a rational model of closed-end fund management that broadly captures the regularities present in the data, something prior researchers have struggled to do. In their influential survey of the literature, Lee et al. (1990) identify the four main empirical regularities exhibited by the closed-end fund discount that they believe any rational model must explain:

1. Closed-end funds are issued at (or above) their NAV, more often than not start trading at a premium to NAV, and then decline.¹⁷
2. On average, closed-end funds trade at a discount relative to their NAV.¹⁸
3. The discount is subject to wide variation over time and across funds.¹⁹
4. At termination, prices converge to NAV.²⁰

The calibration results in Section III show that a simple parsimonious rational model can reproduce all of these features of the data. However, beyond this are the questions of how well the model can *quantitatively* fit the data, and whether it can match other empirical features of the discount. In this section

¹⁷ See Weiss (1989), Peavy (1990), and Weiss Hanley, Lee, and Seguin (1996).

¹⁸ See Pratt (1966) and Zweig (1973).

¹⁹ See Sharpe and Sosin (1975), Thompson (1978), Richards, Fraser, and Groth (1980), Herzfeld (1980), Anderson (1986), and Brauer (1988).

²⁰ See Brauer (1984) and Brickley and Schallheim (1985).

we look at the existing evidence in the literature, and investigate the extent to which it is consistent or inconsistent with the assumptions and predictions of our model. We begin with the empirical evidence supporting the two key assumptions on which our model relies, (1) the existence of managerial ability, and (2) the relation between perceived managerial ability and compensation.

A. Managerial Ability

Despite the high pay of fund managers, and while the idea that some fund managers are better than others does not seem all that controversial, it has proved so difficult for researchers to find convincing evidence of managerial ability that many, such as Jensen (1968) and Carhart (1997), have concluded that it does not exist. However, there are good reasons to expect that evidence of ability would be hard to detect. First, there is the well-known statistical difficulty of directly estimating expected returns: A large amount of data is needed to say conclusively whether one manager can generate expected returns larger than another.²¹ Second, as Berk and Green (2004) point out, even if managers do have ability, and even if large amounts of data are available, ability may not show up in the returns for open-end funds in a competitive market, because the manager captures all of the benefits of that ability.²² Consequently, in the open-end fund sector, the best evidence of the existence of ability is the very strong relation between performance and the flow of funds.

For closed-end funds, while ability should still be hard to detect in the *stock* return of the fund (as we demonstrate in Section III), it should show up in the fund's NAV return. In this context, Dimson and Minio-Paluello (2001) find no evidence of persistence in NAV returns, but Bal and Leger (1996) and Leger (1997) find evidence of performance persistence in U.K. closed-end funds, and Bers and Madura (2000) find similar evidence for U.S. funds. Chen, Rui, and Xu (2001) and Cao and Esman (2002) find evidence of superior performance among managers of Chinese closed-end funds.

Our model further implies that excess NAV returns should be a function of the tenure of the manager, and, in particular, one place ability ought to show up is in the NAV returns of new managers. While Weiss (1989) documents significant *negative* index-adjusted cumulative abnormal NAV returns in the 120 days following a fund's IPO, this is from only a very small sample of funds. More recently, several authors look at fund performance surrounding manager replacement. While infrequent, this is still a much more common event than a fund's IPO. This research finds that new managers exhibit exactly the excess returns predicted by our model. Wermers et al. (2005) study manager replacements in U.S. closed-end funds from 1985 to 2002, and find that, in the year following

²¹ This is shown analytically by Merton (1980).

²² Related to this, the one place where researchers have been able to find evidence of performance persistence in open-end funds is in the worst performing funds, where investors are limited in their ability to withdraw funds by the size of their initial investment (see Brown and Goetzmann (1995), Carhart (1997), Berk and Xu (2005)).

replacement, the new manager generates a cumulative abnormal NAV return, relative to other closed-end funds of the same type, of 1.94%. This number is close to the value predicted by our model in Figure 5, and compares with a cumulative abnormal NAV return of -3.87% in the 2 years prior to replacement. Very similar results are obtained by Rowe and Davidson (2000), who document an excess NAV return of 2% in the year following a manager's replacement. Note that Rowe and Davidson (2000) find no evidence of abnormal stock-price returns following a management change, which is consistent with our model, although Wermers et al. (2005) do find an abnormal stock-price return in the year after the change.

B. Managerial Compensation Contract

An important assumption in our model is that managers become entrenched. The total number of firings in closed-end funds is small: Wermers et al. (2005) document only 260 managerial replacements (which presumably also include voluntary resignations to accept better offers) among 446 closed-end funds in 18 years. This evidence is consistent with bad managers being at least partially entrenched. As Wermers et al. (2005) point out, board directors often have close relationships with the fund's management company, so firing a manager is often difficult. Furthermore, Brauer (1984, 1988) and Brickley and Schallheim (1985) show that although very large gains are associated with managerial turnover in the form of open-ending, this is a very infrequent event, suggesting that large costs must be incurred in order to achieve the turnover.²³

Evidence of rewards for superior performance is harder to find, presumably because many of the rewards for good performance are unobservable. For example, rather than negotiating a pay raise, a manager might choose to spend less time managing the fund and instead spend time either making money elsewhere by managing other funds or increasing his consumption of leisure.²⁴ Nevertheless, Khorana, Wahal, and Zenner (2002) study secondary rights offerings during the period 1988 to 1998, and find that rights offerings are announced, on average, when funds are trading at a premium, and that the premium disappears during the course of the offering. Because a rights offering automatically increases managerial compensation, this evidence is consistent with closed-end fund managers being rewarded for superior performance.

C. Discount Dynamics

Although our model can explain why a fund can be issued at par or at a premium with the expectation that it will later trade at a discount, a widely cited result from Weiss (1989) is that closed-end funds fall to a sizeable discount within only 120 days after their IPO. Our model cannot generate a decrease

²³ See Bradley et al. (2005) for a detailed discussion of these costs.

²⁴ Consistent with this, Wermers et al. (2005) find that the average closed-end manager manages multiple funds.

of this magnitude over such a short period of time for reasonable parameter values. However, more recent research suggests that Weiss's result, which relies on a sample of only 22 U.S. stock fund IPOs between 1985 and 1987, may be very specific to her (small) sample. For example, Levis and Thomas (1995) study 73 U.K. closed-end funds issued between 1984 and 1992, and find that the average time it takes a fund to fall into discount is longer than Weiss finds, and the size of the discount is smaller. Overall, they conclude that the post-IPO performance of closed-end funds is similar to that of industrial IPOs. In the United States, Cherkes et al. (2005) examine a much larger sample of closed-end fund IPOs, over a much longer time period, than Weiss (1989), and find that the average time before a closed-end fund trades at a discount is close to 2 years, a rate of reversion to discount that our model is capable of creating. They find that the average discount across all funds and over the time period of their sample is about 4%, which is within the range of values our model can produce (see Figure 4).

Leaving aside the issue of the short-term post-IPO performance of the fund, there is nevertheless evidence that investors in younger funds earn lower returns than investors in older funds. Each month Cherkes et al. (2005) compare the average stock return of funds that are more than 2 years old to funds of the same type that are less than 2 years old, and find that the older funds outperform the younger funds by about 40 basis points per month. This evidence is difficult to explain within the context of our model, but it does appear to mirror the long-term IPO underperformance evidence.

One thing our model can explain is the wide cross-sectional dispersion in discounts. For example, Weiss (1989) finds that after 24 weeks, while the average discount of the 22 U.S. stock funds she studies is 10%, two of the funds still trade at a premium. Cherkes et al. (2005) find that the cross-sectional standard deviation across funds of the average discount of the fund is 6%. In addition, Thompson (1978, p.159) notes that discounts in excess of "20% are quite frequent, with discounts exceeding 30 and 40% not uncommon." Figure 4 shows that our model is capable of producing this amount of cross-sectional dispersion in discounts.

There are some other results on the dynamics of the discount that our model cannot easily explain. For example, Lee et al. (1991) find that movements in discounts in the same sector are highly correlated, and that the correlation is lower for funds in different sectors. This is confirmed for U.K. funds by Minio-Paluello (1998). Because there is no plausible reason that unexpected changes in managerial quality should be correlated across managers, our model cannot reproduce this feature of the data. However, Cherkes et al. (2005) derive a model of the closed-end fund discount based on the liquidity of different asset classes. Their paper predicts a strong correlation in discounts across funds, indicating that variation in managerial ability cannot be the full story and that liquidity must also play a role in driving discounts. There is no reason why both effects could not be going on simultaneously.

Pontiff (1995) and Thompson (1978) document mean reversion in the closed-end fund discount that can be used to predict abnormal stock returns. Our

model implies a small amount of mean reversion (since there is effectively an upper barrier on the process for ϕ_t), but cannot match the short-term reversals documented in the literature nor can this information be used to generate positive abnormal returns. It is hard to see how an extension of our model could generate this behavior, given the requirement that prices are set competitively, which necessarily implies unpredictable abnormal stock returns. Consequently, this regularity remains a puzzle that needs to be explained.

D. The Discount and NAV Returns

If the discount reflects investors' perception of managers' ability, it ought to be related to returns. In particular, today's discount should be related to past returns (higher returns tell us we have a better manager) and future NAV returns (the better the manager relative to fees charged, the higher the returns he should be able to generate in the future), but not to future stock returns. Furthermore, there should be a contemporaneous correlation between excess NAV returns and changes in the discount (high returns tell us we have a better manager, changing the discount).

As with looking at ability alone, we should expect it to be hard to detect these relations, even if they exist. In addition to the sample size issues mentioned above, there is the further problem that, even if ability is monotonically related to returns, it need not be monotonically related to the discount. As is evident in Figure 1, managers with very high ability are more likely to get a pay raise, and thus the discount may *decrease* as perceived ability rises. Though not explicitly modeled in this paper, Deaves and Krinsky (1994) point out that if poor performance increases the likelihood of firing, the relation between performance and discount may not be monotonic for poor managers either.

With regard to prior returns, Bleaney and Smith (2003) find evidence that, for both U.S. and U.K. funds, the level of the premium is positively correlated with returns over the past 12 to 24 months, and Roenfeldt and Tuttle (1973) find a weak relation between the contemporaneous discount and performance. There is mixed evidence about the relation between NAV returns and the discount. Malkiel (1977) finds no relation between past NAV returns and current discounts. Similarly, Lee et al. (1991) and Pontiff (1995) find an insignificant or negative relation between the premium and the future NAV return. However, Chay (1992), Chay and Trzcinka (1999), and Wu and Xia (2001) find a positive relation between stock fund premiums and future NAV returns. One important difference is that Chay and Trzcinka (1999) do not include bond funds in their sample. Dimson and Minio-Paluello (2001) and Bleaney (2002) are unable to find a relation between the discount and future NAV returns in U.K. data. Bleaney and Smith (2003) find weak evidence that the premium predicts future NAV performance for both U.S. and U.K. funds, but it is not statistically significant. Finally, consistent with the idea that nonmonotonicity might be an important reason why it is hard to find relations between the discount and returns, Wermers et al. (2005) find that the NAV return is positively related to discount returns when managerial replacement is unlikely, but not when replacement is more likely.

V. Conclusions

This paper shows that the primary empirical regularities exhibited by the closed-end fund discount emerge endogenously in a simple, parsimonious rational model driven entirely by the interaction between managerial ability and the manager's labor contract. However, some evidence remains to be explained. For instance, our model requires that investors always earn the competitive market return by investing in the fund. Consequently, the abnormal stock return of the fund should be unpredictable, which is at odds with the findings in Pontiff (1995). Although it is hard to understand how this predictability can persist, the benefit of our model is that it focuses researchers' attention on this particular aspect of the data. Based on the insights from our model, this predictability implies a lack of competition among investors. Precisely what frictions exist that prevent this competition is an interesting question for future research.

Another puzzle is the apparent difference in the stock market performance of old and new funds. Although the evidence in Weiss (1989) about the speed with which discounts appear does not seem to be confirmed in larger samples, there is evidence that younger funds underperform. Why this is the case, and why investors still appear willing to invest in younger funds, remains a mystery and is clearly another question for future research.

Our model also makes several new empirical predictions. For instance, the relation between discounts and returns is highly nonlinear, so it is perhaps not surprising that the linear specifications that have been run to date offer inconclusive results. Instead, the paper suggests that the relation between returns and discounts of very high-ability managers is likely to be inverted, with high returns leading to deeper discounts. An empirical specification based on this insight might yield more definitive results on the relation between discounts and returns.

Other new predictions of the model relate to the impact of managerial turnover. First, new funds (and funds that have recently had a change in management) should exhibit above normal NAV returns and funds with long managerial tenure should exhibit below normal NAV returns. Second, while, consistent with our model, the existing evidence suggests that discounts narrow during proxy fights, as of yet there is no evidence on the behavior of the discount around a *voluntary* management departure or a significant increase in managerial compensation (by, e.g., the manager getting another fund to manage in addition to his existing fund). Our model suggests that only good managers would be willing to leave voluntarily, so that the discount should *increase* (i.e., if the fund is trading at a premium the premium should decrease) around voluntary separations or pay increases.

REFERENCES

- Adams, Andrew T., 2000, Excess volatility and investment trusts, Working paper, University of Edinburgh.
- Anderson, Seth C., 1986, Closed-end funds versus market efficiency, *Journal of Portfolio Management* Fall, 63–67.

- Anderson, Seth C., and Jeffery A. Born, 2002, *Closed-End Fund Pricing* (Kluwer, Boston, MA).
- Arora, Navneet, Nengjiu Ju, and Hui Ou-Yang, 2003, An agency explanation of the closed-end fund puzzles, Working paper, University of Maryland.
- Bal, Yasemin, and Lawrence A. Leger, 1996, The performance of U.K. investment trusts, *Service Industries Journal* 16, 67–81.
- Berk, Jonathan B., and Richard C. Green, 2004, Mutual fund flows and performance in rational markets, *Journal of Political Economy* 112, 1269–1295.
- Berk, Jonathan B., and Jing Xu, 2005, Persistence and fund flows of the worst performing mutual funds, Working paper, U.C. Berkeley.
- Bers, Martina, and Jeff Madura, 2000, The performance persistence of closed-end funds, *Financial Review* 35, 33–52.
- Bleaney, Michael, 2002, Investor sentiment and discounts on closed-end funds, Working paper, University of Nottingham.
- Bleaney, Michael, and R. Todd Smith, 2003, Prior performance and closed-end fund discounts, Working paper, University of Nottingham.
- Boudreaux, Kenneth J., 1973, Discounts and premiums on closed-end mutual funds: A study in valuation, *Journal of Finance* 28, 515–522.
- Bradley, Michael, Alon Brav, Itay Goldstein, and Wei Jiang, 2005, Costly communication, shareholder activism, and limits to arbitrage, Working paper, Duke University.
- Brauer, Gregory A., 1984, Open-ending closed-end funds, *Journal of Financial Economics* 13, 491–507.
- Brauer, Gregory A., 1988, Closed-end fund shares' abnormal returns and the information content of discounts and premiums, *Journal of Finance* 43, 113–128.
- Brickley, James A., Steven Manaster, and James S. Schallheim, 1991, The tax-timing option and the discounts on closed-end investment companies, *Journal of Business* 64, 287–312.
- Brickley, James A., and James S. Schallheim, 1985, Lifting the lid on closed-end investment companies: A case of abnormal returns, *Journal of Financial and Quantitative Analysis* 20, 107–117.
- Brown, Stephen J., and William N. Goetzmann, 1995, Performance persistence, *Journal of Finance* 50, 679–698.
- Cao, Xiaoqing, and Martijn Esman, 2002, Evaluating the performance of Chinese investment funds, Working paper, University of Twente, Netherlands.
- Carhart, Mark M., 1997, On persistence in mutual fund performance, *Journal of Finance* 52, 57–82.
- Chay, Jong-Born, 1992, The pricing of closed-end funds: Discounts and managerial performance, PhD thesis, State University of New York at Buffalo.
- Chay, Jong-Born, and Charles A. Trzcinka, 1999, Managerial performance and the cross-sectional pricing of closed-end funds, *Journal of Financial Economics* 52, 379–408.
- Chen, Gongmeng, Oliver Rui, and Yexiao Xu, 2001, A first look at closed-end funds in China, Working paper, University of Texas at Dallas.
- Chen, Hsiu-Lang, Narasimhan Jegadeesh, and Russ Wermers, 2000, The value of active mutual fund management: An examination of the stockholdings and trades of fund managers, *Journal of Financial and Quantitative Analysis* 35, 343–368.
- Cherkes, Martin, Jacob Sagi, and Richard Stanton, 2005, Liquidity and closed-end funds, Working paper, U.C. Berkeley.
- Chevalier, Judith, and Glenn Ellison, 1999, Career concerns of mutual fund managers, *Quarterly Journal of Economics* 114, 389–432.
- Constantinides, George M., 1983, Capital market equilibrium with personal tax, *Econometrica* 51, 611–636.
- Constantinides, George M., 1984, Optimal stock trading with personal taxes, *Journal of Financial Economics* 13, 65–89.
- Deaves, Richard, and Itzhak Krinsky, 1994, A possible reconciliation of some of the conflicting findings on closed-end fund discounts, *Journal of Business Finance and Accounting* 21, 1047–1057.
- DeLong, J. Bradford, and Andrei Shleifer, 1992, Closed-end fund discounts: A yardstick of small-investor sentiment, *Journal of Portfolio Management* Winter, 46–53.

- Dimson, Elroy, and Carolina Minio-Kozerski, 1999, Closed-end funds: A survey, *Financial Markets, Institutions and Instruments* 8, 1–41.
- Dimson, Elroy, and Carolina Minio-Paluello, 2001, The closed-end fund discount and performance persistence, Working paper, London Business School.
- Dimson, Elroy, and Carolina Minio-Paluello, 2002, *The Closed-End Fund Discount* (Association for Investment Management and Research, Charlottesville, Virginia).
- Ding, Bill, and Russ Wermers, 2004, Mutual fund “stars”: The performance and behavior of U.S. fund managers, Working paper, University of Maryland.
- Ferguson, Robert, and Dean Leistikow, 2004, Closed-end fund discounts and expected investment performance, *Financial Review* 39, 179–202.
- Gemmill, Gordon, and Dylan C. Thomas, 2002, Noise trading, costly arbitrage, and asset prices: Evidence from closed-end funds, *Journal of Finance* 57, 2571–2594.
- Harris, Milton, and Bengt Holmström, 1982, A theory of wage dynamics, *Review of Economic Studies* 49, 315–333.
- Herzfeld, Thomas J., 1980, *The Investor’s Guide to Closed-End Funds* (McGraw-Hill, New York, NY).
- Hu, Fan, Alastair R. Hall, and Campbell R. Harvey, 2000, Promotion or demotion? An empirical investigation of the determinants of top mutual fund manager change, Working paper, Duke University.
- Ingersoll, Jonathan E., Jr., 1976, A theoretical and empirical investigation of the dual purpose funds, *Journal of Financial Economics* 3, 83–123.
- Jensen, Michael C., 1968, The performance of mutual funds in the period 1945–1964, *Journal of Finance* 23, 389–416.
- Khorana, Ajay, 1996, Top management turnover: An empirical investigation of mutual fund managers, *Journal of Financial Economics* 40, 403–427.
- Khorana, Ajay, 2001, Performance changes following top management turnover: Evidence from open-end mutual funds, *Journal of Financial and Quantitative Analysis* 36, 371–393.
- Khorana, Ajay, Sunil Wahal, and Marc Zenner, 2002, Agency conflicts in closed-end funds: The case of rights offerings, *Journal of Financial and Quantitative Analysis* 37, 177–200.
- Kim, Chang-Soo, 1994, Investor tax-trading opportunities and discounts on closed-end mutual funds, *Journal of Financial Research* 17, 65–75.
- Kumar, Raman, and Gregory M. Noronha, 1992, A re-examination of the relationship between closed-end fund discounts and expenses, *Journal of Financial Research* 15, 139–147.
- Lee, Charles M. C., Andrei Shleifer, and Richard H. Thaler, 1990, Anomalies: Closed-end mutual funds, *Journal of Economic Perspectives* 4, 153–164.
- Lee, Charles M. C., Andrei Shleifer, and Richard H. Thaler, 1991, Investor sentiment and the closed-end fund puzzle, *Journal of Finance* 46, 76–110.
- Leger, Lawrence A., 1997, U.K. investment trusts: Performance, timing and selectivity, *Applied Economics Letters* 4, 207–210.
- Levis, Mario, and Dylan Thomas, 1995, Investment trust IPOs: Issuing behavior and price performance. Evidence from the London Stock Exchange, *Journal of Banking and Finance* 19, 1437–1458.
- Liptser, Robert S., and Albert N. Shiryaev, 1978, *Statistics of Random Processes II: Applications*, number 6 in *Applications of Mathematics* (Springer-Verlag, New York, NY).
- Litov, Lubomir, Malcolm Baker, Jessica Wachter, and Jeffrey Wurgler, 2005, Can mutual fund managers pick stocks? Evidence from their trades prior to earnings announcements, Working paper, New York University.
- Malkiel, Burton, 1977, The valuation of closed-end investment company shares, *Journal of Finance* 32, 847–859.
- McIntosh, Willard, Ronald C. Rogers, and C. F. Sirmans, 1994, Stock price and management changes: The case of REITs, *Journal of the American Real Estate and Urban Economics Association* 22, 515–526.
- Merton, Robert C., 1980, On estimating the expected return on the market: An exploratory investigation, *Journal of Financial Economics* 8, 323–361.

- Minio-Paluello, Carolina, 1998, The U.K. closed-end fund discount, PhD dissertation, London Business School.
- Øksendal, Bernt, 2002, *Stochastic Differential Equations: An Introduction with Applications*, 5th ed. (Springer-Verlag, New York, NY).
- Peavy, John W., III, 1990, Returns on initial public offerings of closed-end funds, *Review of Financial Studies* 3, 695–708.
- Pontiff, Jeffrey, 1995, Closed-end fund premia and returns: Implications for financial market equilibrium, *Journal of Financial Economics* 37, 341–370.
- Pontiff, Jeffrey, 1997, Excess volatility and closed-end funds, *American Economic Review* 87, 154–169.
- Pratt, Eugene J., 1966, Myths associated with closed-end investment company discounts, *Financial Analysts Journal* 22, 79–82.
- Richards, R. Malcolm, Donald R. Fraser, and John C. Groth, 1980, Winning strategies for closed-end funds, *Journal of Portfolio Management* Fall, 50–55.
- Roefeldt, Rodney L., and Donald L. Tuttle, 1973, An examination of the discounts and premiums of closed-end investment companies, *Journal of Business Research* 1, 129–140.
- Ross, Stephen A., 2002a, Neoclassical finance, alternative finance and the closed-end fund puzzle, *European Financial Management* 8, 129–137.
- Ross, Stephen A., 2002b, A neoclassical look at behavioral finance; closed-end funds, Princeton Lectures in Finance III.
- Rowe, Wei Wang, and Wallace N. Davidson, III, 2000, Fund manager succession in closed-end mutual funds, *Financial Review* 35, 53–78.
- Seltzer, David, 1989, Closed-end funds: Discounts, premiums and performance, PhD thesis, University of Arizona.
- Sharpe, William F., and Howard B. Sosin, 1975, Closed-end investment companies in the United States: Risk and return, in B. Jacquillat, ed.: *European Finance Association 1974 Proceedings* (North Holland, Amsterdam), pp. 37–63.
- Sias, Richard W., Laura T. Starks, and Seha M. Tiniç, 2001, Is noise trader risk priced? *Journal of Financial Research* 24, 311–329.
- Spiegel, Matthew, 1999, Closed-end fund discounts in a rational agent economy, Working paper, Yale University.
- Thompson, Rex, 1978, The information content of discounts and premiums on closed-end fund shares, *Journal of Financial Economics* 6, 151–186.
- Weiss, Kathleen, 1989, The post-offering price performance of closed-end funds, *Financial Management* Autumn, 57–67.
- Weiss Hanley, Kathleen, Charles M.C. Lee, and Paul J. Seguin, 1996, The marketing of closed-end fund IPOs: Evidence from transaction data, *Journal of Financial Intermediation* 5, 127–159.
- Wermers, Russ, Youchang Wu, and Josef Zechner, 2005, Portfolio performance, discount dynamics, and the turnover of closed-end fund managers, Working paper, University of Maryland.
- Wu, Matthew Qianli, and Yihong Xia, 2001, Co-movement and information content of closed-end fund discounts, Working paper, University of Pennsylvania.
- Zweig, Martin E., 1973, An investor expectations stock price predictive model using closed-end fund premiums, *Journal of Finance* 28, 67–87.