Convex Incentives and Liquidity Premia *

Min Dai † Luis Goncalves-Pinto ‡ Jing Xu § Cheng Yan ¶

March 22, 2022

Abstract

We show that convexities in investors' preferences significantly amplify the effect of transaction costs on liquidity premia. To maximize year-end bonuses, fund managers with negative benchmark-adjusted performance have an incentive to take excessive risks. However, trading costs hinder their ability to do so. Therefore, larger liquidity premia are required to compensate for increased turnover and lower bonuses ensuing from suboptimal risk-taking. These results are robust to alternative sources of convexity such as prospect theory and status concerns. Using data on actively-managed mutual funds, we provide empirical support for the novel predictions of our theoretical model.

Keywords: Mutual Funds, Convex Incentives, Transaction Costs, Liquidity Premia.

JEL Classification: C61, D11, D91, G11.

---

* We appreciate the helpful comments from Adelina Barbalau (discussant), Goncalo Faria, Felix Feng (discussant), Terrence Hendershott, Wenxi Jiang, Peter Kondor, Hong Liu, Dong Lou, Stavros Panageas, Clemens Sialm (discussant), Mikhail Simutin, and Juan Sotes-Paladino, participants at the 2019 WFA Meetings, the 2019 FIRN conference, the 2019 Australasian Finance and Banking conference, and the 2019 FMA Asia-Pacific conference, and seminar participants at Catolica Porto Business School, Chinese University of Hong Kong, Deakin University, Massey University, National University of Singapore, and Renmin University of China.

† Min Dai, National University of Singapore, Department of Mathematics, Risk Management Institute, and NUS Suzhou Research Institute, 10 Lower Kent Ridge Road, S17-08-20, Singapore 119076; Tel. +65 6516-2754; Email: matdm@nus.edu.sg

‡ Luis Goncalves-Pinto (Corresponding Author), University of New South Wales, School of Banking and Finance, Sydney, NSW 2052, Australia; Tel. +612 9385 4291; Email: l.goncalves-pinto@unsw.edu.au

§ Jing Xu, Renmin University of China, School of Finance, Mingde Main Building, 59 Zhongguancun Street, Beijing, China 100872, Email: jing.xu@ruc.edu.cn

¶ Cheng Yan, Essex Business School, University of Essex, Colchester, CO4 3SQ, United Kingdom. Email: cheng.yan.1@cass.city.ac.uk

Convex Incentives and Liquidity Premia

Abstract

We show that convexities in investors' preferences significantly amplify the effect of transaction costs on liquidity premia. To maximize year-end bonuses, fund managers with negative benchmark-adjusted performance have an incentive to take excessive risks. However, trading costs hinder their ability to do so. Therefore, larger liquidity premia are required to compensate for increased turnover and lower bonuses ensuing from suboptimal risk-taking. These results are robust to alternative sources of convexity such as prospect theory and status concerns. Using data on actively-managed mutual funds, we provide empirical support for the novel predictions of our theoretical model.

1 Introduction

It is empirically recognized that liquidity is valuable to investors and that they demand a return premium to compensate for transaction costs. Prior studies find a large difference in expected returns across portfolios sorted on liquidity measures. For instance, Brennan and Subrahmanyam (1996) and Pastor and Stambaugh (2003) find return differences in the order of 6% to 7% per year. However, few existing theories have been able to corroborate the large magnitude of these empirical findings. The seminal work of Constantinides (1986) shows that, even though trading costs significantly reduce the frequency and volume of trading, they have surprisingly little impact on utility. As a result, the return that a marginal investor is willing to exchange for zero trading costs (i.e., the liquidity premium) is found to be an order of magnitude smaller than the transaction cost rate.

There have been prior attempts to reconcile this apparent disconnect between theory and evidence regarding the magnitude of liquidity premia. To the best of our knowledge, the four main references are Jang, Koo, Liu, and Loewenstein (2007), Lynch and Tan (2011), Dai, Li, Liu, and Wang (2016), and Chen, Dai, Goncalves-Pinto, Xu, and Yan (2021). They all argue that the reason for the puzzling disconnect is the assumption of constant investment opportunities used in Constantinides (1986). By extending that seminal model to include time-varying investment opportunities, these studies show that transaction costs can have a much larger effect on liquidity premia. Intuitively, when investment opportunities vary over time, investors are either forced to trade more frequently and pay a heavier transaction cost bill, or they do not trade as frequently but lose utility from highly suboptimal risk exposures.

Our paper complements these prior studies by arguing that convexity in investors' preferences can be an alternative mechanism to increase the magnitude of the effect of transaction costs on liquidity premia, while keeping investment opportunities constant. This is motivated by the observation that convexities are ubiquitous in managerial compensation, behavioural economics, and goal-reaching problems, and that they can significantly affect the decision-making process. We study this question both theoretically and empirically, by focusing on the risk-taking incentives induced by the year-end bonus payments that are prevalent in the compensation contracts of mutual fund managers. However, our main results are robust to other non-concave incentives from the behavioural economics literature, such as those derived from prospect theory (Kahneman and Tversky (1979)) and status concerns (Lee, Zapatero, and Giga (2018)).

The mutual fund industry provides an ideal setting to examine how convex incentives affect the magnitude of liquidity premia, for four main reasons. First, this industry has grown rapidly in the last two decades.[1] Hence, it is reasonable to assume that mutual fund managers as a group have become the marginal investor in many stocks.[2] Second, the compensation contracts of fund managers typically include a convex component in the form of year-end bonus payments (Ma, Tang, and Gomez (2019)). They reward good performance but do not penalize bad performance, giving incentives to take excessive risks (Lee, Trzcinka, and Venkatesan (2019)).[3] Moreover, these contracts are typically incomplete, giving fund managers the residual control over the assets in their portfolios (Hart (2017)).[4] Third, fund managers that are prone to engage in excessive risk-taking also tend to hold more illiquid stocks in their portfolios compared to other funds (Huang, Sialm, and Zhang (2011)), implying that risk-taking induced by convex incentives is especially likely to interact with stock trading costs. Fourth, data on mutual fund characteristics and their portfolio holdings

[1] According to the Investment Company Institute, the total net assets of US-registered investment companies was \$22.5 trillion at the end of 2017, compared to \$13 trillion at the end of 2007, which represents a 73% increase. The growth was even more dramatic one decade prior, with an increase of about 177% in total net assets from a value of \$4.7 trillion at the end of 1997.

[2] For instance, Boguth and Simutin (2018) show that the tightness of leverage constraints in mutual funds, as captured by their demand for high-beta stocks, is a priced risk factor in the cross-section of stock returns. This is consistent with mutual funds being the marginal investors.

[3] Lee, Trzcinka, and Venkatesan (2019) show that risk-taking of fund managers is motivated more by this type of compensation structure than by a tournament to capture flows. There is a long strand of empirical research on the tournament incentives derived from the convex flow-performance relation, starting with Brown, Harlow, and Starks (1996), and Chevalier and Ellison (1997). However, some studies have shown that the tournament effect disappears when using different measures, different methodologies, and different data frequencies. This is why we focus our study on the year-end bonus component that is typical in fund manager's compensation contracts. For instance, Busse (2001) shows that the evidence on mutual fund tournaments disappears when using daily data. Kempf, Ruenzi, and Thiele (2009) show that the risk-taking incentive of fund managers is contingent on the state of the economy and on employment risk. Schwarz (2012) argues that mean reversion in the volatility of fund returns mechanically generates the tournament effect. Spiegel and Zhang (2013) challenge the economic foundation for the existence of a convex relation between performance and subsequent fund flows, and they show that it is in fact linear when properly estimated.

[4] Fund managers have numerous ways to change the riskiness of their portfolios, such as switching between equity and cash, or switching between low and high beta stocks, among others. This suggests that it may be very difficult, if not impossible, to specify all the contingencies in their contracts to deter excessive risk-taking, which deems these contracts incomplete.

is readily available, allowing us to directly test the empirical predictions of our theoretical model.

To investigate the effect of convex incentives on liquidity premia, we introduce proportional transaction costs in the model of Basak, Pavlova, and Shapiro (2007).[5] The model considers a risk-averse fund manager whose performance is measured relative to an external benchmark. The manager can invest in a risk-free bond and a benchmark stock that are both perfectly liquid, and she can invest in a non-benchmark stock which is subject to transaction costs and provides some unspanned risk relative to the benchmark stock.[6] If the manager is able to beat the benchmark by year-end, she is rewarded with a bonus from her compensation contract. However, if she does poorly relative to the benchmark, there is no penalty.

In the presence of year-end bonuses, the fund manager's effective risk appetite is a function of the performance of the fund relative to the benchmark. The manager's willingness to accept gambles is higher when the fund is underperforming the benchmark, because in this region the impact of the gamble on the marginal value of the manager's compensation is also larger. Thus, it is optimal for the fund manager to distort the portfolio away from the benchmark and increase its tracking-error volatility when underperforming to improve the odds of finishing ahead by year-end and capture the bonus.[7] This is commonly referred to as the risk-shifting incentive. In addition, the fund manager also has a lock-in incentive. That is, when the fund's portfolio performance overtakes that of the benchmark, it is optimal for the fund manager to lock-in this relative advantage by tracking the benchmark closely.

An immediate implication of this risk-shifting versus lock-in policy is that it requires high portfolio turnover to build up and undo the potentially large portfolio deviations from the

[5] This is a non-trivial task, because the introduction of transaction costs renders the market incomplete, in which case the martingale technique is no longer applicable.

[6] We include two stocks, liquid and illiquid, to match the empirical evidence in Huang, Sialm, and Zhang (2011) that risk-takers tend to hold more illiquid stocks. However, in Section 3.3 we corroborate our findings using alternative models with a single stock, like in Constantinides (1986).

[7] The manager first deviates from the benchmark by reducing the holdings of liquid benchmark stock in exchange for risk-free bond. This is because the perfectly liquid stock is a more efficient risk-taking tool. However, after hitting the position limits on the liquid stock, the fund manager needs to load up on the illiquid non-benchmark stock. Mutual funds typically have tight position limits on their portfolio holdings. In Section 3.2.3 we discuss the implications associated with the relaxation of such position limits.

benchmark.[8] Such high turnover requirements imply heavy trading cost bills to be incurred by fund managers. We simulate the optimal policy under our baseline model and find that the expected discounted value of the trading cost payments in the presence of convex incentives is more than sixty times larger than in the case without such incentives.

In addition to the direct costs associated with high portfolio turnover, the introduction of trading costs also reduces the effectiveness of risk-shifting as a strategy to capture year-end bonuses, because adjusting portfolio positions frequently is expensive. Simulation results show that when we increase the trading cost rate from 0% to 1%, the manager's risk-adjusted bonus decreases by about 10%.[9]

As marginal investors, fund managers demand high liquidity premia to compensate for (i) the heavy trading cost payments associated with increased portfolio turnover, and (ii) the reduction in total compensation due to suboptimal risk-shifting. We decompose the liquidity premia into these two parts and show that, as the transaction cost rate increases, the second component becomes the main contributor to the amplification effect in our model.

Following Constantinides (1986), we focus our analysis on the ratio between the model-implied liquidity premium and the transaction cost rate (LPTC). In his seminal paper, Constantinides finds that the model-implied LPTC ratio is only 0.07, concluding that the liquidity premium is an order of magnitude smaller than the trading cost rate and is therefore negligible.[10]

We use a collar function to model year-end bonuses, and match the moments of the

[8] This is consistent with the high portfolio turnover observed in practice for mutual funds. According to the Investment Company Institute, FactBook 2018, the asset-weighted average portfolio turnover rate of equity mutual funds, for the period 1984-2017, is around 57%. The implication that benchmark-linked incentives lead to higher turnover is not novel. This result has been obtained by Cuoco and Kaniel (2011) and Sotes-Paladino and Zapatero (2019). However, their models do not account for the effects of transaction costs.

[9] The manager's risk-adjusted bonus is defined as the minimum amount of additional assets under management (AUM) that the manager requires for foregoing the bonus. This measure is a function of managerial risk aversion. The 10% reduction includes the mechanical effect of trading cost payments on relative performance. If we assume that trading costs are waived upon trading, to prevent such mechanical effect, the manager's risk-adjusted bonus decreases by 4.61% when increasing the transaction rate from 0% to 1%. For fair comparability, when we waive the trading cost payments, the optimal policy that is followed is a policy that includes the no-trading region generated by those trading costs.

[10] These results are generated using a proportional transaction cost rate of 1% for both purchases and sales, which is equivalent to a 2% round-trip charge. We use this as baseline in the rest of the paper, unless stated otherwise.

bonus distribution to the empirical estimates provided in Ma, Tang, and Gomez (2019). We find that our model-implied LPTC ratio is around 1.175, implying that trading costs are a first-order determinant of liquidity premia in our setting.[11] When we remove the convexity from our model, the LPTC ratio drops to 0.02, which is even lower than in Constantinides (1986).[12] For a trading cost rate of 1%, the liquidity premium generated by our model with bonuses is sixty nine times larger compared to an equivalent model without bonuses.

The main empirical prediction of our model is that, the stronger the risk-taking incentives induced by the convexity in the compensation structure, the larger the liquidity premia. We use a sample of U.S. domestic actively-managed equity mutual funds and their portfolio holdings to provide support for this prediction. To the best of our knowledge, this is the first time in the literature that this relation has been established and tested.

We start by constructing fund-level variables that proxy for convex incentives in the mutual fund industry. Specifically, we assume that fund managers who have engaged in risk-shifting in the past to be more likely to continue to do it in the future. Therefore, we use the eight risk-shifting measures proposed by Huang, Sialm, and Zhang (2011). These measures capture the various ways fund managers can use to change the riskiness of their portfolios, such as changing the portfolio composition between equity holdings and cash holdings, and within equity holdings switching between low beta and high beta stocks, or changing the idiosyncratic risk of the portfolio, by increasing the tracking-error volatility to their benchmarks, or increasing portfolio concentration in certain industries or styles.

We aggregate each of the eight fund-level proxies across all funds holding a given stock, using quarterly share holdings as weights. Then, for each stock, we aggregate across the eight proxies by averaging the cross-sectional percentile ranks or by using the first principal

[11] We assume that the fund manager receives a bonus only if her portfolio outperforms the external benchmark. Our results remain qualitatively similar when we use any of the alternative specifications proposed in Basak, Pavlova, and Shapiro (2007). These additional results are available from the authors upon request.

[12] In contrast to Constantinides (1986), the investor in our model does not derive utility from intermediate consumption, does not have an infinite investment horizon, and has access to a liquid stock which can be correlated with the illiquid stock. It is important to consider these differences when comparing and contrasting our results with those in Constantinides (1986). For tractability, we attribute the liquidity premium entirely to the illiquid stock, but in equilibrium it could potentially be split between the two stocks. In Section 3.3 we examine alternative settings with a single stock in which the liquidity premium can only be attributed to that stock, and the results are qualitatively similar to our baseline results.

component. This results in a single stock-level proxy that captures the convex incentives of the mutual funds that hold that stock.[13] We then examine how this proxy for stock-level convex incentives affects the relation between transaction costs and future stock returns.

We regress excess returns on lagged trading costs interacted with convex incentives. In our list of control variables we include stock characteristics such as size, turnover, and ownership by active mutual funds, as well as the factor loadings from a variety of factor models. We use the effective trading cost estimates and the testing methodology of Hasbrouck (2009). We confirm that effective trading costs are strongly related to future stock returns. More importantly, we find that the interaction of trading costs with convex incentives is significant at the 1% level and is economically large: for a 1% increase in effective trading costs, portfolios of stocks with strong risk-shifting incentives require more than 1% higher excess return per month. These results are consistent with our theoretical model, that risk-shifting incentives in the mutual fund industry are first-order determinants of the liquidity premia of stocks, keeping size, turnover, ownership, and other stock characteristics, constant.

In our theoretical model, the two sources of utility losses for fund managers are the trading cost charges associated with portfolio rebalancing, and the suboptimal risk-shifting that results in reduced compensation. In the data, we find that the average turnover of the portfolios of stocks held by funds with strong risk-shifting incentives is about double that in the weak incentive group. The difference is even more significant for stocks with low betas, in which case the turnover goes from two to two and a half times larger. However, for stocks with high betas, turnover is only about 36% larger in the strong incentive group. This is consistent with the argument in Boguth and Simutin (2018) that low beta stocks provide lower leverage, and risk-shifting with such stocks requires larger trades. We also show in the empirical analysis that the relation between trading costs and turnover is more negative in the high incentive group, suggesting that fund managers are more likely to shift risk using liquid stocks, requiring additional future return to use illiquid stocks in their risk-shifting

[13] In a previous version of this paper we have used each of the eight proxies individually from the first-level aggregation, i.e., the aggregation across all funds holding a given stock. The results are qualitatively similar using each proxy compared to those reported in this paper using the second-level aggregation, i.e., the aggregation of the eight stock-level proxies into a single one. The results are available from the authors upon request.

activities.

The remainder of this paper unfolds as follows. Section 2 presents the theoretical framework. Section 3 describes the numerical analysis of the optimal investment policy, the magnitude of the liquidity premia, and their sensitivity to changes in parameter values. Section 4 provides empirical evidence to support the novel predictions of the theoretical model. Section 5 concludes. We relegate to the Appendix all the technical issues, additional results, and details on the construction of the empirical variables.

2 Theoretical Framework

We introduce proportional transaction costs in the model of Basak, Pavlova, and Shapiro (2007). Time is continuous and there exist three assets in the economy: a risk-free bond (S_{0t}) and a benchmark stock (S_{1t}) that are perfectly liquid, and a non-benchmark stock (S_{2t}) that is subject to transaction costs. We assume that S_{0t} grows at a constant rate r and S_{it} evolves according to the following process:

$$dS_{it} = \alpha_i S_{it} dt + \sigma_i S_{it} dW_{it} \tag{1}$$

for $i = 1, 2$, where the two standard Brownian motions W_{1t} and W_{2t}, defined on a filtered complete probability space $(\Omega, \mathcal{F}, P)$, have constant correlation $\rho \in [-1, 1]$. The expected returns (α_i) and volatilities (σ_i) are assumed to be constant, like in Constantinides (1986).

We consider a mutual fund manager who invests in these three assets. She can buy the illiquid non-benchmark stock for the price $(1 + \lambda)S_{2t}$, and she can sell it for the price $(1 - \mu)S_{2t}$, where $\lambda \geq 0$ and $0 \leq \mu < 1$ represent the proportional transaction cost rates for purchases and sales, respectively. Let X_t denote the dollar amount invested in the bond and in the liquid stock, and let Y_t denote the dollar amount invested in the illiquid stock. We then have the following sub-wealth processes for liquid and illiquid portfolio holdings:

$$\begin{aligned} dX_t &= (rX_t + \xi_t(\alpha_1 - r))dt + \sigma_1 \xi_t dW_{1t} - (1 + \lambda)dL_t + (1 - \mu)dM_t & (2) \\ dY_t &= \alpha_2 Y_t dt + \sigma_2 Y_t dW_{2t} + dL_t - dM_t & (3) \end{aligned}$$

where ξ_t in (2) is the dollar amount invested in the liquid holdings, and L_t and M_t are non-decreasing processes which denote the cumulative amounts of purchases and sales of the illiquid stock, respectively, which have initial values $L_{0-} = M_{0-} = 0$.

We assume the fund manager's compensation at the end of the investment period (T) consists of two parts: a management fee proportional to the value of assets under management (AUM), and a bonus that is determined by the fund's performance relative to an external benchmark. We assume that the benchmark consists of liquid stock and bond, and is continuously rebalanced to maintain a $\beta/(1-\beta)$ stock-bond ratio. Therefore, its value Z_t evolves according to the following process:

$$dZ_t = (r + \beta(\alpha_1 - r))\, Z_t dt + \beta\sigma_1 Z_t dW_{1t} \tag{4}$$

Let $f \equiv f(R_T^f - R_T^b)$ represent the ratio of the manager's bonus to the management fee. Then, the manager's total compensation at time T equals:

$$k\left(1 + f(R_T^f - R_T^b)\right)(X_T + Y_T) \tag{5}$$

where k is the management fee rate, and $R_T^f = \ln\frac{X_T+Y_T}{X_0+Y_0}$ and $R_T^b = \ln\frac{Z_T}{Z_0}$ are the continuously compounded gross returns over the period $(0, T)$ for the fund and the benchmark, respectively. Since only the growth rate of the fund's portfolio over the benchmark matters for the calculation of the bonus, we can set $Z_0 = X_0 + Y_0$ without loss of generality.

We consider a bonus function of the collar-type:

$$f(R_T^f - R_T^b) = \begin{cases} f_L & \text{if } R_T^f - R_T^b < \theta_L \\ f_L + \psi(R_T^f - R_T^b - \theta_L) & \text{if } \theta_L \le R_T^f - R_T^b < \theta_H \\ f_H \equiv f_L + \psi(\theta_H - \theta_L) & \text{if } R_T^f - R_T^b \ge \theta_H \end{cases} \tag{6}$$

where $f_L \ge 0$, $\psi = (f_H - f_L)/(\theta_H - \theta_L) > 0$, and $\theta_L < \theta_H$. This bonus function exhibits a local convexity around the lower threshold θ_L, then increases linearly until it reaches the upper threshold θ_H, above which it returns to a flat position.[14]

[14] We have examined alternative bonus functions. The results are qualitatively similar to those we obtain

We assume that the manager's objective is to maximize the expected utility she derives from the amount of total compensation at time T, which is equivalent to:

$$V(0, X_0, Y_0, Z_0) = \max_{\Theta_{[0,T]}} \; E\left[\frac{((1+f)(X_T + Y_T))^{1-\gamma}}{1-\gamma}\right] \tag{7}$$

where $\gamma > 0$ and $\gamma \neq 1$ is the manager's risk aversion coefficient, and $\Theta_{[0,T]} \equiv \{(\xi_s, L_s, M_s) : 0 \leq s \leq T\}$ denotes the manager's investment policy over the period $[0, T]$.[15]

We impose borrowing and short-selling constraints in our model, and the fund manager's admissible investment policies are such that $0 \leq \xi_t \leq X_t$ and $Y_t \geq 0$ for all t. In practice, mutual funds are typically subject to prohibitions against borrowing and short-selling (Almazan, Brown, Carlson, and Chapman (2004)), limits on tracking-error and cash holdings (Simutin (2014)), and style-drift restrictions such as SEC Rule 35(d)-1.

We include two stocks, liquid and illiquid, to match the empirical evidence in Huang, Sialm, and Zhang (2011) that risk-shifters tend to hold more illiquid stocks. However, in Section 3.3 we corroborate our findings using alternative models with a single stock, like in Constantinides (1986).

3 Model Implications

This section presents a numerical analysis of the fund manager's optimal policy, and the liquidity premia implied by the model. We solve the manager's problem numerically, since a closed-form solution is not available.[16]

In our baseline model, we use the following parameter values. The expected return and volatility of the liquid benchmark stock are chosen to match the average annual return and

using this collar specification.

[15] Like in Dai, Jin, and Liu (2011), our fund manager derives utility from the gross assets rather than the liquidated assets of the fund. This avoids trading strategies that lead to liquidation at T and helps focus on the effect of transaction costs on interim trading. In other words, when we compare the cases with or without transaction costs, the amount of bonuses at T are identical conditional on the same level of relative performance. This provides a conservative estimate of the effect of transaction costs. As expected, when the fund manager derives utility from liquidated wealth, the results are strictly stronger than the currently reported.

[16] Appendix A describes the solution method and the numerical procedure, and Appendix B discusses equilibrium implications.

volatility, over the period 1950-2017, of the S&P 500 index, with $\alpha_1 = 9\%$ and $\sigma_1 = 14\%$. The expected return and volatility of the illiquid non-benchmark stock are chosen to match the average annual return and volatility, over the same period, of the value-weighted portfolio formed by the stocks in the lowest decile of market capitalization. This gives $\alpha_2 = 19\%$ and $\sigma_2 = 24\%$, and a return correlation of $\rho = 0.53$ with the S&P 500 index. We assume that the benchmark is fully invested in liquid stock, i.e., $\beta = 1$. The baseline transaction cost rate is set at 1%, i.e., $\lambda = \mu = 1\%$. The risk-free rate is estimated from the average T-bill return over the period 1950-2017, which is $r = 4\%$. We set the fund manager's risk aversion level at $\gamma = 5$, and the investment horizon at one calendar year ($T = 1$).

We match the parameter values in the bonus function (6) to the empirical estimates of Lee, Trzcinka, and Venkatesan (2019) and Ma, Tang, and Gomez (2019): $\theta_L = 0.01$, $\theta_H = 0.15$, $f_L = 0$, and $f_H = 1.5$. These imply that, when the performance of the fund's portfolio in excess of the benchmark is below 1% annually, the fund manager receives no bonus. If instead the fund's portfolio outperforms the benchmark by at least 15% annually, the fund manager receives a bonus that amounts to 150% of her management fee. When the relative performance is between 1% and 15%, the amount of bonus increases linearly with respect to relative performance.

3.1 Optimal Stock Allocations

In our model, at a given time t, the optimal policy is a function of the performance of the fund relative to the benchmark, i.e., $\eta = \ln \frac{X_t+Y_t}{Z_t}$. Figure 1 illustrates the optimal allocations in benchmark and non-benchmark stocks at mid-year (i.e., $t = T/2$). Panel A shows the case without transaction costs (i.e., $\lambda = \mu = 0\%$), and Panel B shows the case with a trading cost rate of 1% (i.e., $\lambda = \mu = 1\%$). Both panels show that the optimal policy entails two types of actions: (i) to deviate from the benchmark when underperforming ($\eta < 0$), and (ii) to lock-in the relative advantage when outperforming ($\eta > 0$). The fund manager achieves (i) by first reducing the holdings of liquid benchmark stock in exchange for risk-free bond until hitting the position limits. This is because the liquid stock is a more efficient risk-shifting tool. After hitting the position limits on liquid stock, the fund manager starts overweighting the illiquid non-benchmark stock in the portfolio. The objective of such portfolio deviations is to increase the likelihood that the fund will outperform the benchmark by the terminal date, when the bonus is calculated. This is called the risk-shifting range in Basak, Pavlova, and Shapiro (2007). The fund manager achieves (ii) by unwinding the entire position on the illiquid non-benchmark stock and the risk-free bond to track the benchmark very closely. This is especially the case when the relative performance approaches the upper threshold θ_H.[17] For extreme values of the relative performance, either positive or negative, the fund manager follows the normal policy, which is similar to the policy in Merton (1969) for the case without trading costs, as represented by the horizontal sections of the policy in Panel A.

[Insert Figure 1 about here]

Panel B presents the case with transaction costs. Like in other portfolio choice models with transaction costs (e.g., Constantinides (1986), Davis and Norman (1990), Liu and Loewenstein (2002), Liu (2004), Chellathurai and Draviam (2007), Jang, Koo, Liu, and Loewenstein (2007), Dai, Li, Liu, and Wang (2016)), the optimal policy is characterized by a no-trading region for the non-benchmark stock. The no-trading region is delimited by buy and sell boundaries, which also exhibit large swings when relative performance η switches between positive and negative values. For extreme values of relative performance, either positive or negative, the fund manager chooses a constant range of risk exposure to the non-benchmark stock (see the horizontal sections of Panel B), which is independent of relative performance, like in Constantinides (1986).

This optimal policy implies the following trading behaviour for the fund manager. When the portfolio weight for the non-benchmark stock is pushed to the sell region (the area above the sell boundary), then the fund manager sells this stock instantaneously to push its portfolio

[17] In our baseline case, we impose the typical position limits that are pervasive in practice, such as leverage and short-sale constraints (Almazan, Brown, Carlson, and Chapman (2004)). We highlight that, given the baseline parameter values we use in our model, risk-shifting is not the only reason why the fund manager holds the illiquid non-benchmark stock. This stock can also provide diversification benefits, given its imperfect correlation with the benchmark stock in our baseline case, and it can be held for the risk premium it offers. This is different from the models in Basak, Pavlova, and Shapiro (2007) and Dai, Goncalves-Pinto, and Xu (2019), where the non-benchmark stock is such that it only carries idiosyncratic risk. We cannot use such setup, because this would make it impossible to estimate the model-implied liquidity premia in Section 3.2, which is the central focus of our paper.

weight back to the sell boundary. If the portfolio weight on the non-benchmark stock falls into the buy region (the area below the buy boundary), the manager buys instantaneously to push it back to the buy boundary. If the portfolio weight on the non-benchmark stock is between the buy and sell boundaries, the fund manager is better off not trading this stock, because the improvement in risk exposure is more than offset by the costs incurred with trading.

The trading boundaries vary dramatically with η. It is very likely that these boundaries will be hit due to changes in relative performance. The frequency and volume of trading in the non-benchmark stock are then expected to increase as a result. This is in contrast with the optimal policy in Constantinides (1986), which is independent of relative performance.

Table 1 reports some statistics for the manager's optimal investment policy, to understand the effects of convex incentives on trading. We present three panels for different trading cost rates. In the first row of each panel, we present the case in which convex incentives are absent from the fund manager's problem. The amount of trading is tiny in this case, which is consistent with the findings of Constantinides (1986) that even small transaction cost charges can significantly reduce the volume and frequency of trading. For instance, in the absence of convex incentives, and for a transaction cost rate of 1% (Panel B), the discounted value of the transaction costs ($PVTC$) paid during the investment period is only 0.001% of the fund's initial AUM. In addition, the fund manager trades only 0.36 times per year on average, assuming at most two trades per business day.

[Insert Table 1 about here]

In the presence of convex incentives, the amount of trading increases dramatically, which in turn leads to a heavy trading cost bill. For example, in Panel B, when the transaction cost rate is 1%, in the presence of bonuses the $PVTC$ rises to 0.481% of the fund's initial AUM, which is about 480 times higher than in the case without bonuses. This is indicative that the effect of transaction costs on the manager's derived utility is much larger in a model with convex incentives.

3.2 Liquidity Premia

In this section, we define liquidity premium, and decompose its sources into two parts: (i) the part that is directly driven by the trading cost charges associated with stock turnover, and (ii) the part that is due to suboptimal risk-shifting, as portfolio rebalancing is difficult due to trading costs. We also examine the sensitivity of liquidity premia to changes in the values of the input parameters.

3.2.1 Definition and Decomposition

Constantinides (1986) defines liquidity premium as the maximum expected return an investor is willing to forgo in exchange for zero transaction costs. We define the liquidity premium in our model in a similar way. It is the quantity δ that the fund manager is willing to subtract from the expected return of the illiquid non-benchmark stock to reduce the transaction cost rates μ and λ to zero. This is equivalent to the following indifference condition:

$$V(0, X_0, Y_0, Z_0; \mu, \lambda, \alpha_2) = V(0, x, y, z; 0, 0, \alpha_2 - \delta) \tag{8}$$

where the value function V is defined in (7). The equation above shows that we measure the liquidity premium as of $t = 0$, and that the premium is attributed entirely to the illiquid stock. We also assume that the manager builds her initial portfolio at zero cost, and relaxing this assumption would only strengthen our results.

The manager's utility is affected by the presence of transaction costs through two channels: (i) the payment of transaction costs directly reduces the value of AUM, and (ii) the fund manager is forced to have a risk exposure that is suboptimal, compared to the case without transaction costs. To differentiate these two effects, we follow Dai, Li, Liu, and Wang (2016). We redefine the indifference condition as follows:

$$V^A(0, X_0, Y_0, Z_0; \mu, \lambda, \alpha_2) = V(0, x, y, z; 0, 0, \alpha_2 - \delta^0)$$

where V^A denotes the value function when (i) the manager implements the optimal trading strategy assuming positive transaction costs and (ii) the transaction costs are waived upon trading. Therefore, δ^0 now captures the liquidity premium that can be attributed solely to the suboptimal risk exposure. This is contrast with δ in equation (8), which measures the total liquidity premium. Thus, δ^0/δ measures the fraction of the liquidity premium that is due to portfolio displacement, relative to the total premium.

Table 2 reports the characteristics of the optimal trading policy (at time $t = 0$) and the liquidity premia for various levels of transaction costs. To facilitate comparison, we first present in Panel A the liquidity premia when the bonus is absent (δ_c). This case is analogous to Constantinides (1986), except that we assume a shorter investment horizon, no intermediate consumption, and the presence of a second stock that is perfectly liquid. We find that the liquidity premia are very small in this case. For example, when the transaction cost rate is 1% (2%), the liquidity premium to transaction cost (LPTC) ratio (i.e., $\delta_c/(\lambda+\mu)$) is 0.017 (0.009). If the fund manager is the marginal investor in this stock, her trading on this stock is expected to command a negligible liquidity premium.

[Insert Table 2 about here]

However, in the presence of bonuses (Panel B) the liquidity premia can be very large. For instance, assuming a 1% (2%) transaction cost rate, the liquidity premium (δ) is 2.35% (3.58%), which translates to an LPTC ratio of about 1.175 (0.895). The last row in Panel B shows that the liquidity premia in the presence of bonuses (δ) can be 49 to 136 times larger in magnitude compared to the case without bonuses (δ_c).

Panel B also reports the maximum and minimum values for the buy and sell boundaries of the manager's optimal investment policy. For instance, if the trading cost rate is 1%, the buy boundary goes from a minimum of $B_*(0,\eta) = 0.00$ to a maximum of $B^*(0,\eta) = 0.55$. Figure 1 shows that the minimum is reached when the fund's portfolio is outperforming the benchmark and it is optimal to lock-in the relative advantage by replicating the benchmark closely. The maximum value is reached somewhere within the risk-shifting range. The corresponding max and min values for the selling boundary are $S^*(0,\eta) = 0.74$ and $S_*(0,\eta) = 0.05$, respectively.

Table 2 also shows the proportion of the liquidity premium that is due to direct payments and to portfolio distortions. For example, for a 1% transaction cost rate, the ration δ^0/δ equals 94.68% in the absence of bonuses (Panel A), and 50.77% in the presence of bonuses

(Panel B). This implies that direct payments contribute more to liquidity premia in the presence of bonuses, which is consistent with the results reported in Table 1.[18]

Panel B also shows that the costs due to portfolio displacement are can be substantial in the presence of bonuses. This is because when transactions are costly to execute, the manager's ability to influence her future relative performance is significantly hampered. Consequently, the manager's ability to capture bonuses is weakened, which results in substantial costs of portfolio displacement. To verify this intuition, we perform Monte-Carlo simulations to calculate the manager's risk-adjusted bonus (RAB), which is defined as the minimum amount of additional AUM that the manager requires for waiving her year-end bonus. In addition, we consider the case in which the manager follows her optimal policy under transaction costs, but these costs are waived. This helps isolate the effect of suboptimality in portfolio composition from the direct trading cost payments. Figure 2 shows that the suboptimality in risk-shifting created by the presence of transaction costs substantially reduces RAB. For example, when the transaction cost rate increases from 0% to 1%, RAB decreases by 4.61%. If trading costs were not waiver, this figure would be larger.

[Insert Figure 2 about here]

Our two-stock model matches the empirical findings in Huang, Sialm, and Zhang (2011), that risk-shifters in the mutual fund industry hold more illiquid stocks. However, our main result that convexity in investor's preferences generates much higher liquidity premia also holds for a single stock. In Section 3.3 we consider other non-concave incentives derived from prospect theory and status concerns, in which we use a single stock.

The convexity of the bonus component is the driver of the effects in our model. If we were to consider a pure benchmarking effect (i.e., where the manager derives concave utility from the wealth-to-benchmark value ratio, like in Binsbergen, Brandt, and Koijen (2008)), this would not result in larger liquidity premia.

[18] The relative contribution of direct payments versus portfolio distortions, in the case without bonuses (i.e., δ_c^0/δ_c), is larger in our model compared to Dai, Li, Liu, and Wang (2016) for two main reasons. First, we consider a relatively short investment horizon of one year, and the investor does not trade when near maturity. As a result, the cost due to portfolio displacement in our model is larger. In contrast, they assume a long horizon, and this effect is weaker in their model. Second, our investor can also trade a perfectly liquid stock, which is positively correlated with the illiquid stock. This is expected to reduce her demand for trading the illiquid stock in our model.

3.2.2 Comparative Statics

We report comparative statics analyses for the liquidity premia with regards to a battery of model parameters. Table 3 reports the results of such analyses, some of which we discuss below.

[Insert Table 3 about here]

Return correlation: The correlation between the returns of benchmark and non-benchmark stocks is an important determinant of the effectiveness of risk-shifting. Intuitively, stocks with low correlation with the benchmark are better tools for risk-shifting purposes. If the correlation increases, then the benefit from risk-shifting decreases and the fund manager trades less of the illiquid stock. This is then expected to reduce the liquidity premia. Table 3 suggests that this is indeed the case. For example, when the return correlation is increased by 10% from its baseline value of 0.53, the LPTC ratio decreases from 1.175 to 1.138.

Riskiness of the benchmark portfolio: The riskiness of the benchmark portfolio also affects the manager's risk-shifting incentives. To outperform a riskier benchmark the manager needs to load more on the illiquid non-benchmark stock, as it provides some risk unspanned by the benchmark. This in turn leads to larger liquidity premia. Table 3 shows the changes in the LPTC ratio against changes in β, which measures the riskiness of the benchmark portfolio. It shows that the liquidity premia increase with β. For example, for a 80-20 stock-bond benchmark portfolio (i.e., $\beta = 0.8$), the LPTC ratio is 1.168, which increases to 1.175 when the benchmark portfolio is fully invested in liquid stock (i.e., $\beta = 1$), as in the baseline case.

Convexity of the bonus-performance relationship: The convexity of the bonus function is the main driver of the fund manager's risk-shifting behaviour. Keeping all else constant, increasing the convexity of this function creates stronger risk-shifting incentives, and the liquidity premia increase as a result. We can increase the convexity of the bonus-performance relation by either shrinking the range of relative performance over which the function exhibits an upward-sloping pattern, i.e., $[\theta_L, \theta_H]$, or by increasing the reward for outperformance, i.e., f_H. Table 3 shows that the LPTC ratio increases when we either reduce θ_H or increase f_H, implying that the liquidity premia increase with the convexity of the bonus function. This is the main prediction that we test in our empirical analysis of Section 4.

3.2.3 Discussion on Position Limits

In our baseline model we exogenously impose the typical position limits that exist in practice for mutual funds, such as limits on leverage and short-selling (see Almazan, Brown, Carlson, and Chapman (2004)). It could be argued that our results are driven by such position limits. This issue is especially critical for the liquid benchmark stock, which is a better risk-shifting tool due to its high liquidity. In our baseline, the fund manager is restricted from borrowing to buy the benchmark stock on margin, and is restricted from shorting that stock as well.

The manager first deviates from the benchmark by reducing the holdings of liquid benchmark stock in exchange for risk-free bond. This is because the liquid stock is a more efficient risk-shifting tool. However, after hitting the position limits on the liquid stock, the fund manager loads on the illiquid non-benchmark stock.

Therefore, we should expect that, when position limits are relaxed, the demand for illiquid stock will decrease, and the impact of trading costs on the derived utility of the fund manager will be weaker.

Figure 3 shows what happens when we relax the position limits on the liquid benchmark stock. We continue to fix the position limit on the illiquid non-benchmark stock to be within the interval $[0, 1]$.[19] However, we relax the position limit on the liquid benchmark stock to vary within the interval $[-a, 1 + a]$, where a ranges from 0 to 2. The results show that the LPTC ratio decreases slightly for looser position limits on the liquid benchmark stock.

[Insert Figure 3 about here]

We conclude that position limits on the liquid benchmark stock are not the main driver of our results.

[19] If we relax the position limit of the illiquid non-benchmark stock, the results only become stronger.

3.3 Alternative Sources of Convexity: Reference-Dependent Utility

In our baseline model, the utility function of the fund manager is globally concave, and the manager's incentives stem from the convexity of her compensation contract, namely the existence of year-end bonuses. However, in the economics literature more generally, some utility specifications are designed to exhibit local convexity more directly. Many reference-dependent utilities examined in the behavioural economics literature have this property. In this section, we show that convex incentives arising from some reference-dependent utilities can also generate high liquidity premia. Thus, the effect that we document is not restricted to the specific model that we examine in our baseline framework.

We assume, similar to Constantinides (1986), the investor only trades in one risk-free bond and one risky stock, but with a major difference in terms of the utility specification. The investor's objective is to choose her investment policies to maximize

$$E[U(W_T; R)] \tag{9}$$

where W_T is the investor's gross wealth level at time T, and R is a reference point which determines the investor's utility level.[20]

We consider two popular forms of reference-dependent utility. First, the prospect theory model proposed in Kahneman and Tversky (1979), is widely used in behavioural finance studies. This utility function has the following specification:

$$U(W; R) = \begin{cases} (W-R)^p & \text{if } W \geq R \\ -c(R-W)^q & \text{if } W < R \end{cases} \tag{10}$$

where $0 < p, q < 1$ and $c > 0$. This function is defined on deviations of wealth (W) from the reference point (R). It is concave for gains $(W \geq R)$ and convex for losses $(W < R)$. It results in an S-shaped function with an inflexion at the reference point (R). Investors with

[20] Like in our main model, we assume the investor derives utility from gross wealth level at time T, to exclude the mechanical effect of transaction cost paymnets on wealth resulting from liquidation.

such utility function will be risk-averse relative to gains and risk-loving relative to losses, which leads to the well-known behavioural bias coined as disposition, where investors realize gains quickly but ride losses.

Second, the aspiration utility examined in Diecidue and van de Ven (2008) and Lee, Zapatero, and Giga (2018), which has the following specification:

$$U(W;R) = \begin{cases} \frac{W^p}{p} & \text{if } W < R \\ c_1\frac{W^p}{p} + c_2 & \text{if } W \geq R \end{cases} \tag{11}$$

where $0 < p < 1$, $c_1 \geq 1$ and $c_2 \geq 0$. This specification captures the idea that, besides normal consumption, the investor also cares about her status, which is revealed through the consumption of non-divisible goods, such as a luxury car or an apartment. Thus, the investor's utility will jump when her wealth reaches the level beyond which she is able to consume the non-divisible good.

In the following analysis, the default parameter values that we use are as follows: the risk-free rate is $r = 0.04$, the expected return of the risky stock is $\alpha = 0.1$, the return volatility of the risky stock is $\sigma = 0.3$, and the investor's investment horizon is $T = 1$ year. The parameters in function (10) are calibrated to the estimates of Kahneman and Tversky (1979), as follows: $p = q = 0.88$ and $c = 2.25$. The reference point is set at the investor's initial wealth level W_0. The parameter values in function (11) are as follows: $p = 0.5$, $c_1 = 1.2$, $c_2 = 0$, and the reference point is set at $R = 1.2W_0$.[21]

Figure 4 plots the investor's optimal allocation to the risky stock, as a function of her wealth level. This figure is generated in the absence of transaction costs. In both cases, we find that the optimal allocation in the risky stock changes dramatically around the reference point R. This implies that the investor is likely to frequently adjust her stock allocations in response to fluctuations in wealth. In particular, when the investor's wealth level is below the reference point (i.e., lies in the convex region of her utility function), she increases the stock allocation, implying a strong risk-shifting incentive. Therefore, similar to the intuition developed in our baseline model, it is expected that the presence of transaction costs will be

[21] For brevity, we omit the details of the mathematical model. They are available from the authors upon request.

particularly burdensome to an investor who has incentives to distort take excessive risks to increase the odds of overcoming the reference point.

[Insert Figure 4 about here]

Figure 5 plots the LPTC ratio against the transaction cost rate. The liquidity premium is calculated at the start of the investment period. The figure shows that, the liquidity premium can be substantial when the investor exhibits either prospect theory utility or aspiration utility. For example, with a transaction cost rate of 1%, the LPTC ratio is above 0.7 in the prospect theory utility case, and it is above 1.0 in the aspiration utility case. Figure 5 also shows the LPTC ratio when we increase the convexity of these utility functions. The convexity of the prospect theory utility is increased by reducing q from 0.88 to 0.6, and the convexity of the aspiration utility is increased by reducing R from $1.2W_0$ to $1.1W_0$. The results suggest that greater convexity of the utility functions result in larger liquidity premia. This is consistent with our baseline model of Section 2.

[Insert Figure 5 about here]

4 Empirical Analysis

In this section, we provide empirical evidence to support the main implications of our theoretical model. We show that trading costs have a positive relation with future stock returns, especially for stocks held by mutual funds with strong risk-shifting incentives. We also show that mutual fund risk-shifting is positively associated with portfolio turnover. However, increased turnover is not the only explanation for the positive relation between trading costs and future returns. To the best of our knowledge, the existing empirical evidence has not yet verified any of these novel implications of our model.

4.1 Sample

To test our model, we require information on mutual funds, their quarterly share holdings, and characteristics of the stocks they hold, especially a measure of stock trading costs. We obtain mutual fund returns, investment objectives, fees, total net assets (TNA), and other fund characteristics from the Center for Research in Security Prices (CRSP) Survivor-Bias-Free Mutual Fund Database. We use the Wharton Research Data Services (WRDS) MFLINKS file to merge this database with Thomson Financial Mutual Fund Holdings (TFMFH), which contains information on the number of shares of stock held by each mutual fund, and provides identifiers for fund families (Wermers (2000)).[22]

The version of the MFLINKS file that we use in this analysis includes the mapping of mutual funds to their portfolio holdings from 1980 to 2016. This sets the upper limit of our sample coverage to the end of 2017, because in our analysis we use information of mutual fund holdings with a one-year lag. However, we limit our sample to start in 2004, which is the year when the SEC imposed a new regulation requiring more frequent (quarterly) portfolio disclosures by mutual funds. This choice of sample coverage from 2004 to 2017 can only bias our results against our main hypothesis. This is because such increase in the frequency of portfolio disclosures by mutual funds had a positive effect on stock liquidity, as documented in Agarwal, Mullally, Tang, and Yang (2015).[23] In addition, in 2004 the CRSP mutual fund database switched its data provider from Morningstar to Lipper. The overlap between these two versions of the database is imperfect. By focusing on the period after 2004 we guarantee consistency in the reported data.

We restrict our analysis to diversified domestic actively-managed equity mutual funds.[24] Fund-level variables are aggregations across all the share classes. For instance, the Total Net Assets (TNA) of the fund is the sum of the TNAs of all its share classes.

Reported fund objectives do not always accurately characterize a fund. Thus, following Kacperczyk, Sialm, and Zheng (2008) and Glode (2011), we exclude funds that hold on average less than 80% of their net assets in equity. We also exclude funds with TNA lower than $5 million, because Elton, Gruber, and Blake (2001) show that the returns of such

[22] TFMFH only reports large portfolio holdings, i.e., with a dollar value of at least $100 million. This is unlikely to bias our results in favour of our main hypothesis, because the reported holdings are more likely to be larger and more liquid stocks in the CRSP universe.

[23] More generally, stock market liquidity has improved significantly in the more recent period. Ben-Rephael, Kadan, and Wohl (2015) show that the liquidity premium is indiscernible from zero in the past two decades. However, we stress that the effective trading cost estimates in our sample are not negligible, and we are focusing our analysis on the ratio between the liquidity premium and the trading cost rate (LPTC), and this ratio can still be large even if the liquidity premium is small.

[24] We exclude international, balanced, sector, bond, money market, and index funds.

small funds tend to be biased upwardly in the CRSP database.[25]

We extract stock-level information, such as prices, returns, volume, and shares outstanding, from the CRSP Stock Database. We only use common stocks (i.e., share codes 10 and 11). We then merge this data with the mutual fund holdings from TFMFH. The average number of stocks per month in our final sample is 1,877, and the average number of funds per quarter is 1,776.

We have computed the number of stock shares held by the active mutual funds in our final sample, as a percentage of the total shares outstanding. The time-series of the cross-sectional average of this ownership ratio is reported in Figure 6. The graph shows that the largest increase in ownership by active funds occurs in 2004, with a 24% increase in average quarterly ownership compared to 2003. This result also supports our focus on the period after 2004, because it is more likely that these funds will be marginal investors in these stocks during this period.

[Insert Figure 6 about here]

4.2 Main Variables

The main prediction of our theoretical model is that risk-shifting in the mutual fund industry is an important factor driving the relation between transaction costs and future stock returns, i.e., the liquidity premium. To test this prediction, we need to identify mutual fund managers that are more (or less) prone to engage in risk-shifting, and we need a measure of trading costs.

Starting with the latter, we use the effective trading cost estimates of Hasbrouck (2009) and his testing method. The data is available from the author's website, but it ends in 2009.[26] We extend his dataset until 2017, following his estimation method. Hasbrouck (2009) performs separate analyses for NASDAQ, AMEX, and NYSE stocks. However, we focus our analysis on NASDAQ stocks, like in Eleswarapu (1997). This is because the average

[25] We require 36 months of return history for our analysis, which mitigates the issue of incubation bias discussed in Evans (2010).

[26] http://people.stern.nyu.edu/jhasbrou/Research/GibbsCurrent/gibbsCurrentIndex.html

number of stocks in NASDAQ is three to four times larger than in the other exchanges. We use extra dimensions for sorting stocks into portfolios, and focusing on NASDAQ allows us to obtain around 35 stocks per portfolio. Using the stocks from the other exchanges would result in a much lower number.

After merging all the datasets of the previous section with this dataset on effective trading costs, our final sample has 1,760 stocks on average per cross-section.

Regarding the proxies for risk-shifting incentives, we adopt the study by Huang, Sialm, and Zhang (2011). They measure risk-shifting of a mutual fund by comparing the risk of their most recently disclosed portfolio with the risk of the fund's returns in the past. If their most recent portfolio is riskier than the fund's returns in the past, then the fund is taking additional risks. They use eight alternative measures of risk, including total volatility, idiosyncratic volatility, tracking-error volatility, etc. These measures capture the various ways that fund managers can use to change the riskiness of their portfolios, such as changing the portfolio composition between equity holdings and cash holdings, and within equity holdings switching between low beta and high beta stocks, or changing the idiosyncratic risk of the portfolio, by increasing the tracking-error volatility to their benchmarks, or increasing portfolio concentration in certain industries or styles. We provide a detailed description of each of these eight measures in Appendix C.

We start by calculating each of these fund-level measures for every mutual fund in our sample. Next, we aggregate each measure across all funds holding a given stock, using their quarterly share holdings as weights. This gives us one stock-level measure for each of the eight risk-shifting proxies. We format these measures such that large values correspond to stronger incentives by the holders of the stock to engage in risk-shifting.

We then create a composite measure at the stock-level. We use two methods to accomplish this. First, we use the average percentile rank across the eight measures. In December of the prior year, we assign percentile ranks to each stock-level measure for the whole cross-section of stocks. Then, we compute the average percentile rank across the eight measures for a given stock. We denote this composite measure as APR.

The second composite measure uses principal component analysis. In December of the prior year, we compute the principal components across the eight stock-level measures, and use the first principal component obtained from that analysis. We denote this composite measure as FPC.[27]

4.3 Portfolio Statistics and Empirical Specification

In this section, we test how the risk-shifting incentives (RSI henceforth) of the stockholders affect the relation between effective trading costs and future stock returns. We proxy for RSI using either one of the composite measures, APR or FPC, as described above.

For this test, we adopt the methodology used in Hasbrouck (2009). This allows us to directly compare our results with his, as we use the same effective trading cost estimates.

We start by forming portfolios based on a sequential three-way sort. First, we rank stocks based on the average RSI in the prior year, and we form two portfolios using the median as cutoff. Second, within each of these two portfolios, we sort stocks into quintiles based on the beta of the stocks, which is estimated using a market model over a 36-month lookback window. Third, within each of the ten portfolios formed with RSI and beta, we sort stocks into quintiles based on the effective trading cost measure. In total, this three-way sort results in 50 portfolios, and the number of stocks per portfolio is on average 35. We study the monthly returns of these 50 portfolios over a 14-year period from 2004 to 2017, which means our final sample includes a total of 8,400 portfolio-month observations.

In our theoretical model, we show that convex incentives increase stock turnover. This is confirmed in the data. Table 4 shows that, when we compute RSI using APR in Panel A (FPC in Panel B), the average turnover of the stocks in the above-median RSI portfolios is about 66% (67%) higher than the stocks in the below-median group.

[Insert Table 4 about here]

The difference in turnover is particularly large for the portfolios of stocks with low betas. In Panel A, the stocks in the bottom quintile of beta (i.e., beta rank equal to "Low") exhibit

[27] These aggregations are useful as we do not need to report multiple iterations of the same tests for each of the eight stock-level proxies. We report only one test per aggregated measure. However, we have also used the individual stock-level measures instead of the composites in our empirical analysis, and the results are qualitatively similar. These results are available from the authors upon request.

an average turnover that is about 143% larger in the above-median RSI group, compared with the below-median RSI group. The turnover of the stocks in the top quintile of beta (i.e., beta rank equal to "High") is only about 36% larger in the above-median group. The difference is similar in Panel B. These results are consistent with the argument in Boguth and Simutin (2018) that low beta stocks provide lower leverage, and taking risks with such stocks requires larger trades.

Next, we examine the relation between trading costs and future stock returns, and how this relation is affected by convex incentives in the mutual fund industry. We regress future returns (in excess of the risk-free rate) of the 50 portfolios on the lagged effective trading cost measure of Hasbrouck (2009) (averaged across all stocks in each portfolio) and its interaction with an indicator variable for above-median RSI, which we denote as $DummyRSI$. The dependent variable is the equal-weighted average monthly return across the stocks within each portfolio. The portfolio sorting variables (i.e., RSI, beta, and trading cost) are all measured in the year prior to the year of the portfolio returns.

The full empirical specification for this test is as follows:

$$\begin{aligned} R_{i,t+1} = & \gamma_0 + \gamma_c c_{i,t} + \gamma_d DummyRSI_{i,t} + \gamma_{cd} \ (c_{i,t} \times DummyRSI_{i,t}) \\ & + \gamma_t Turnover_{i,t} + \gamma_{ct} \ (c_{i,t} \times Turnover_{i,t}) \\ & + \gamma_l LRMC_{i,t} + + \gamma_o Ownership_{i,t} \\ & + \gamma_m \beta_i^m + \gamma_s \beta_i^{smb} + \gamma_h \beta_i^{hml} + \epsilon \end{aligned} \tag{12}$$

where β^m, β^{smb}, and β^{hml} are the unconditional betas obtained from the Fama and French (1993) three-factor model, estimated over the entire sample period for each portfolio. $LRMC$ is the log relative market capitalization (i.e., the average median-adjusted market capitalization for the stocks in each portfolio). The betas and $LRMC$ were also used in Hasbrouck (2009) as controls. We add to his baseline specification three additional controls, as well as interactions with the effective trading cost measure c. Specifically, we include controls for (i) $Turnover$, which is the average ratio of trading volume to shares outstanding across the stocks in each portfolio, (ii) $Ownership$, the ratio of the number of shares held by active mutual funds to the total number of shares outstanding, averaged across stocks in the portfolio, and (iii) $DummyRSI$, an indicator function that equals one for above-median RSI portfolios, and zero otherwise. We interact the indicator function with the effective trading cost (i.e., $c \times DummyRSI$), to assess the impact of RSI on the relation between trading costs and future stock returns. Following Hasbrouck (2009), we estimate this regression using the Generalized Method of Moments (GMM).

4.4 Regression Results

We start by confirming the findings in Hasbrouck (2009) that effective trading costs are strongly related to future stock returns. In columns (1) and (6) of Table 5 we report positive and significant coefficients on the effective trading cost variable (c).

[Insert Table 5 about here]

The difference between Panels A and B is one variable used in the sequential triple sorting to create the 50 portfolios used in the analysis. In Panel A we use APR, and in Panel B we use FPC, as our stock-level proxy for risk-shifting incentives. This explains the slight differences in the regression coefficients for the trading cost variable across the two panels. The tests in columns (1) and (6) control only for the variables used in Hasbrouck (2009).

In columns (2) and (7) we control for ownership by the active funds in our sample. This guarantees that we are comparing stocks held to the same extent by the funds in our sample.

In columns (3) and (8) we include the interaction between trading costs and the indicator $DummyRSI$. We find the coefficient on the interaction to be statistically significant and economically large. For instance, in column (3) of Panel A the interaction coefficient suggests that, for a 1% increase in effective trading costs, portfolios in the above-median RSI group require 1.43% higher return per month, compared to those in the below-median RSI group. Given that the average return in the below-median group is 1.74% per month (i.e., coefficient of c, when $DummyRSI = 0$), this is equivalent to 82% larger return for the above-median RSI group. The results are even stronger in Panel B where FPC is used instead of APR.

We have shown in Table 4 that stocks held by funds with above-median RSI exhibit higher turnover, which is consistent with our theoretical predictions. However, Lee and Swaminathan (2000) have shown that stock return momentum is strongly affected by past trading volume. Specifically, they find that stocks with high (low) past turnover ratios earn lower (higher) future returns. This explains the negative coefficient on $DummyRSI$ in columns (3) and (8), and suggests that we should control for stock turnover in our tests to rule out this potential alternative driver of returns.

In columns (4) and (9), we replace $DummyRSI$ with $Turnover$.[28] As expected, the coefficient on the interaction between trading costs and turnover is positive and significant. That is, the higher the turnover, the stronger the effect of trading costs on future returns.

In columns (5) and (10), we report the results for the full specification, in which the main variable of interest is the indicator $DummyRSI$ and its interaction with trading costs, but we control for turnover and its interaction with trading costs as well. This is to address the potential concern that stocks in the above-median RSI groups could be mechanically associated with high turnover.

We conclude that these additional controls do not explain away the amplification effect that our RSI measures have on the relation between trading costs and future stock returns. For instance, in column (5), the coefficient on the interaction $c \times DummyRSI$ has a value of 1.07, which is statistically significant at the 1% level. This implies that, keeping size, ownership, and turnover constant, the average monthly return required by above-median RSI portfolios is nearly double the return required by the below-median RSI group (i.e., the coefficient on c, when $DummyRSI = 0$), which is an economically large effect.

In Panel C, we include additional betas from alternative factor models. In columns (1)-(2) and (5)-(6) we include the unconditional momentum beta (*UMD Beta*) from the four-factor model of Carhart (1997) (FF4). In columns (3)-(4) and (7)-(8) we include the unconditional betas of profitability (*RMW Beta*) and investment (*CMA Beta*) from the five-factor model of Fama and French (2015) (FF5). Our main results remain qualitatively unchanged.

Overall, these results suggest that convex incentives in the mutual fund industry are important determinants of the liquidity premia of stocks. This is evidence consistent with the main implication of our theoretical model.

[28] We use turnover instead of trading volume, because the former is a much less skewed variable.

4.5 Convex Incentives and Turnover

In the decomposition exercise of Section 3.2.1, we show that a significant portion of liquidity premium is generated by the suboptimal risk-shifting imposed by the presence of trading costs. This suggests that the increased turnover associated with risk-shifting may not always be the main driver of liquidity premia. To provide evidence consistent with this, we estimate an additional model in which we use portfolio turnover as the dependent variable.

We expect to find the relation between trading costs and turnover to be more negative for stocks held by funds with stronger risk-shifting incentives. This would be consistent with risk-shifting leading to higher turnover for more liquid stocks. In other words, to use illiquid stocks in their risk-shifting activities, fund managers should require additional compensation in the form of larger future stock returns.

Table 6 reports the results of a regression model with the following specification:

$$\begin{aligned} Turnover_{i,t+1} = & \gamma_0 + \gamma_c\, c_{i,t} + \gamma_d \;\; DummyRSI_{i,t} + \gamma_{cd} \;\; \left(c_{i,t} \times DummyRSI_{i,t}\right) \\ & + \gamma_s Ln(Size)_{i,t} + \gamma_a Alpha_{i,t} + \gamma_b Beta_{i,t} + \gamma_v IdioVol_{i,t} \\ & + \gamma_d DivYield_{i,t} + \gamma_o Ownership_{i,t} + \epsilon_{i,t} \end{aligned} \tag{13}$$

where $Turnover$ is the average turnover of the stocks in each of the 50 portfolios, created following the sequential sorting procedure described above. The indicator *DummyRSI* and the trading cost measure (c) are defined as in the previous section. We control for the characteristics examined in Lo and Wang (2000): (i) the natural log of a stock's market capitalization, averaged across all stocks in a portfolio (*Ln(Size)*), (ii) the intercept coefficient from the time-series regression of a stock's return on the value-weighted market return, averaged across all stocks in a portfolio ($Alpha$), (iii) the slope coefficient from the time-series regression of a stock's return on the value-weighted market return, averaged across all stocks in a portfolio ($Beta$), (iv) the residual standard deviation of the time-series regression of a stock's return on the value-weighted market return, averaged across all stocks in a portfolio (*IdioVol*), and (v) the average dividend yield of the stocks in each portfolio (*DivYield*). These five characteristics have been considered important determinants of stock turnover. In addition to these five variables, we also control for ownership by active mutual funds, as in the previous section.

[Insert Table 6 about here]

Table 6 shows that, the variables examined in Lo and Wang (2000) are strongly related with future turnover. But the main coefficient of interest is on the interaction ($c \times DummyRSI$), which compares the effect of trading costs on turnover between stocks with high and low convex incentives. The interaction is negative and statistically significant at the 1% level, with t-statistics ranging from -7.82 to -10.19. This suggests that trading costs have a significantly more negative effect on stock turnover for stocks held by funds with stronger risk-shifting incentives. These results are consistent with our theoretical results. They imply that the amplification effect found in Table 4 cannot be fully explained by an increase in turnover, but instead by the suboptimal risk exposure induced by the presence of trading costs.[29]

5 Conclusion

In the mutual fund industry, contracts are typically incomplete and include option-like components, which induce fund managers to engage in excessive risk-taking. We use this industry as a laboratory to study how convex incentives created by compensation contracts affect the liquidity premia of stocks.

Theoretically, we show that the optimal response to the convexities embedded in their contracts is to deviate from the benchmark when underperforming, and lock-in when outperforming. In the presence of transaction costs, the high portfolio turnover implicit in the optimal investment strategy generates a heavy trading cost bill. Moreover, trading costs also make it more difficult for fund managers to adjust their portfolio positions, which makes risk-shifting less effective at capturing year-end bonuses.

[29] These results on turnover can also shed some light on the apparently inconsistent findings in Hasbrouck (2009). This prior work shows that, the coefficient on effective cost is too large (> 1) to be consistent with a simple trading story. However, we show that, under convex incentives, suboptimal risk exposure can play a more important role than trading expenses in explaining liquidity premia, and this can help explain the large magnitude of the coefficients found in Hasbrouck (2009).

We find that fund managers demand high liquidity premia to compensate for two effects: (i) the direct trading cost payments associated with increased portfolio turnover, and (ii) the bonuses that are lost due to the portfolio rebalancing difficulties imposed by trading costs.

Empirically, we show that the interaction between trading costs and convex incentives is significant and of first-order importance at explaining future returns. This effect is not driven exclusively by increased turnover, implying that it must also be driven by portfolio distortions that are hard to undo because of trading costs, hampering fund managers' ability to capture year-end bonuses. Therefore, fund managers require higher future returns to be willing to take risks with illiquid stocks.

These results suggest that convex incentives in the mutual fund industry are important determinants of the liquidity premia of stocks. To the best of our knowledge, this is the first time such results have been reported in the literature.

References

Agarwal, V., K. Mullally, Y. Tang, and B. Yang, 2015, "Mandatory Portfolio Disclosure, Stock Liquidity, and Mutual Fund Performance," *Journal of Finance*, 70, 2733–2776.

Almazan, A., K. C. Brown, M. Carlson, and D. A. Chapman, 2004, "Why Constrain Your Mutual-Fund Manager?," *Journal of Financial Economics*, 73, 289–321.

Basak, S., A. Pavlova, and A. Shapiro, 2007, "Optimal Asset Allocation and Risk Shifting in Money Management," *Review of Financial Studies*, 20, 1583–1621.

Ben-Rephael, A., O. Kadan, and A. Wohl, 2015, "The Diminishing Liquidity Premium," *Journal of Financial and Quantitative Analysis*, 50, 197–229.

Binsbergen, J. H. V., M. W. Brandt, and R. S. J. Koijen, 2008, "Optimal Decentralized Investment Management," *Journal of Finance*, 63, 1849–1895.

Boguth, O., and M. Simutin, 2018, "Leverage Constraints and Asset Prices: Insights from Mutual Fund Risk Taking," *Journal of Financial Economics*, 127, 325–341.

Brennan, M. J., and A. Subrahmanyam, 1996, "Market microstructure and asset pricing: On the compensation for illiquidity in stock returns," *Journal of Financial Economics*, 41, 441–464.

Brown, K. C., W. V. Harlow, and L. T. Starks, 1996, "Of Tournaments and Temptations: An Analysis of Managerial Incentives in the Mutual Fund Industry," *Journal of Finance*, 51, 85–110.

Buss, A., and B. Dumas, 2019, "The Dynamic Properties of Financial-Market Equilibrium with Trading Fees," *Journal of Finance*, 74, 795–844.

Busse, J., 2001, "Another Look at Mutual Fund Tournaments," *Journal of Financial and Quantitative Analysis*, 36, 53–73.

Carhart, M. M., 1997, "On Persistence in Mutual Fund Performance," *Journal of Finance*, 52, 57–82.

Chellathurai, T., and T. Draviam, 2007, "Dynamic Portfolio Selection with Fixed and/or Proportional Transaction Costs using Non-Singular Stochastic Optimal Control Theory," *Journal of Economic Dynamics and Control*, 31, 2168–2195.

Chen, Y., M. Dai, L. Goncalves-Pinto, J. Xu, and C. Yan, 2021, "Incomplete Information and the Liquidity Premium Puzzle," *Management Science*, 67, 5703–5729.

Chevalier, J., and G. Ellison, 1997, "Risk Taking by Mutual Funds as a Response to Incentives," *Journal of Political Economy*, 105, 1167–1200.

Constantinides, G. M., 1986, "Capital Market Equilibrium with Transaction Costs," *Journal of Political Economy*, 94, 842–862.

Cuoco, D., and R. Kaniel, 2011, "Equilibrium Prices in the Presence of Delegated Portfolio Management," *Journal of Financial Economics*, 101, 264–296.

Dai, M., L. Goncalves-Pinto, and J. Xu, 2019, "How Does Illiquidity Affect Delegated Portfolio Choice?," *Journal of Financial and Quantitative Analysis*, 54, 539–585.

Dai, M., H. Jin, and H. Liu, 2011, "Illiquidity, Position Limits, and Optimal Investment for Mutual Funds," *Journal of Economic Theory*, 146, 1598–1630.

Dai, M., P. Li, H. Liu, and Y. Wang, 2016, "Portfolio Choice with Market Closure and Implications for Liquidity Premia," *Management Science*, 62, 368–386.

Dai, M., and Y. Zhong, 2010, "Penalty Methods for Continuous-Time Portfolio Selection with Proportional Transaction Costs," *Journal of Computational Finance*, 13, 1–31.

Davis, M. H., and A. R. Norman, 1990, "Portfolio Selection with Transaction Costs," *Mathematics of Operations Research*, 15, 676–713.

Diecidue, E., and J. van de Ven, 2008, "Probability of Success and Failure, and Expected Utility," *International Economic Review*, 49, 683–700.

Eleswarapu, V., 1997, "Cost of Transacting and Expected Returns in the Nasdaq Market," *Journal of Finance*, 52, 2113–2127.

Elton, E. J., M. J. Gruber, and C. R. Blake, 2001, "A First Look at the Accuracy of the CRSP Mutual Fund Database and a Comparison of the CRSP and Morningstar Mutual Fund Databases," *Journal of Finance*, 56, 2415–2430.

Evans, R. B., 2010, "Mutual Fund Incubation," *Journal of Finance*, 65, 1581–1611.

Fama, E., and K. French, 2015, "A Five-Factor Asset Pricing Model," *Journal of Financial Economics*, 116, 1–22.

Fama, E. F., and K. R. French, 1993, "Common Risk Factors in the Returns of Stocks and Bonds," *Journal of Financial Economics*, 33, 3–56.

Forsyth, P. A., and K. R. Vetzal, 2002, "Quadratic Convergence of a Penalty Method for Valuing American Options," *SIAM Journal on Scientific Computation*, 23, 2095–2122.

Glode, V., 2011, “Why Mutual Funds “Underperform”,” *Journal of Financial Economics*, 99, 546–559.

Goncalves-Pinto, L., J. Sotes-Paladino, and J. Xu, 2018, “The Invisible Hand of Internal Markets in Mutual Fund Families,” *Journal of Banking and Finance*, 89, 105–124.

Hart, O., 2017, “Incomplete Contracts and Control,” *American Economic Review*, 107, 1731–1752.

Hasbrouck, J., 2009, “Trading Costs and Returns for U.S. Equities: Estimating Effective Costs from Daily Data,” *Journal of Finance*, 64, 1445–1477.

Huang, J., C. Sialm, and H. Zhang, 2011, “Risk Shifting and Mutual Fund Performance,” *Review of Financial Studies*, 24, 2575–2616.

Jang, B.-G., H. K. Koo, H. Liu, and M. Loewenstein, 2007, “Liquidity Premia and Transaction Costs,” *Journal of Finance*, 62, 2330–2366.

Kacperczyk, M., C. Sialm, and L. Zheng, 2008, “Unobserved Actions of Mutual Funds,” *Review of Financial Studies*, 21, 2379–2416.

Kahneman, D., and A. Tversky, 1979, “Prospect Theory: An Analysis of Decision under Risk,” *Econometrica*, 47, 263–291.

Kempf, A., S. Ruenzi, and T. Thiele, 2009, “Employment Risk, Compensation Incentives, and Managerial Risk Taking: Evidence from the Mutual Fund Industry,” *Journal of Financial Economics*, 92, 92–108.

Lee, C., and B. Swaminathan, 2000, “Price Momentum and Trading Volume,” *Journal of Finance*, 55, 2017–2069.

Lee, J. H., C. Trzcinka, and S. Venkatesan, 2019, “Do Portfolio Manager Contracts Contract Portfolio Management?,” *Journal of Finance*, 74, 2543–2577.

Lee, S., F. Zapatero, and A. Giga, 2018, “Rolling the Skewed Die: Economic Foundations of the Demand for Skewness,” Working Paper, University of Southern California.

Liu, H., 2004, “Optimal Consumption and Investment with Transaction Costs and Multiple Risky Assets,” *Journal of Finance*, 59, 289–338.

Liu, H., and M. Loewenstein, 2002, “Optimal Portfolio Selection with Transaction Costs and Finite Horizons,” *Review of Financial Studies*, 15, 805–835.

Lo, A. W., and J. Wang, 2000, "Trading Volume: Definitions, Data Analysis, and Implications of Portfolio Theory," *Review of Financial Studies*, 13, 257–300.

Lynch, A. W., and S. Tan, 2011, "Explaining the Magnitude of Liquidity Premia: The Roles of Return Predictability, Wealth Shocks, and State-Dependent Transaction Costs," *Journal of Finance*, 66, 1329–1368.

Ma, L., Y. Tang, and J.-P. Gomez, 2019, "Portfolio Manager Compensation in the U.S. Mutual Fund Industry," *Journal of Finance*, 74, 587–638.

Merton, R. C., 1969, "Lifetime Portfolio Selection Under Uncertainty: The Continuous-Time Case," *Review of Economics and Statistics*, 51, 307–318.

Pastor, L., and R. F. Stambaugh, 2003, "Liquidity risk and expected stock returns," *Journal of Political Economy*, 111, 642–685.

Schwarz, C. G., 2012, "Mutual Fund Tournaments: The Sorting Bias and New Evidence," *Review of Financial Studies*, 25, 913–936.

Shreve, S. E., and H. M. Soner, 1994, "Optimal Investment and Consumption with Transaction Costs," *Annals of Applied Probability*, 4, 609–692.

Simutin, M., 2014, "Cash Holdings and Mutual Fund Performance," *Review of Finance*, 14, 1425–1464.

Sotes-Paladino, J., and F. Zapatero, 2019, "Riding the Bubble with Convex Incentives," *Review of Financial Studies*, 32, 1589–1622.

Spiegel, M., and H. Zhang, 2013, "Mutual Fund Risk and Market Share Adjusted Fund Flows," *Journal of Financial Economics*, 108, 506–528.

Wermers, R., 2000, "Mutual Fund Performance: An Empirical Decomposition into Stock-Picking Talent, Style, Transaction Costs, and Expenses," *Journal of Finance*, 55, 1655–1703.

Appendix

The content of this Appendix is as follows. Appendix A describes the method used to solve our baseline problem. Appendix B offers a short discussion on equilibrium implications. Appendix C describes the construction of the empirical proxies used in Section 4.

Appendix A. Solution Method to the Fund Manager's Problem

We solve the fund manager's problem using dynamic programming. For this purpose, we define the value function for $0 \leq t \leq T$ as follows,

$$V(t,x,y,z) = \max_{\Theta_{[t,T]}} \; E\left[\frac{[(1+f)(X_T+Y_T)]^{1-\gamma}}{1-\gamma} | X_t = x, Y_t = y, Z_t = z\right] \tag{14}$$

Under regularity conditions, $V(t,x,y,z)$ must satisfy the following Hamilton-Jacobi-Bellman (HJB) equation (cf. Shreve and Soner (1994)):

$$\min\big\{ -\partial_t V - \mathcal{L}V, \;\; \partial_y V - (1-\mu)\partial_x V, \;\; (1+\lambda)\partial_x V - \partial_y V\big\} = 0 \tag{15}$$

with terminal condition

$$V(T,x,y,z) = \frac{1}{1-\gamma}\left[\left(1+f\left(\ln\frac{x+y}{z}\right)\right)(x+y)\right]^{1-\gamma} \tag{16}$$

in the solution domain $\Omega = \{(t,x,y,z) : 0 \leq t \leq T, x \geq 0, y \geq 0, z \geq 0\}$, where ∂ represents the partial derivative operator, and the differential operator $\mathcal{L}$ is given by:

$$\begin{aligned}\mathcal{L}V \;\; &= \;\; rx\partial_x V + \alpha_2 y\partial_y V + (r+\beta(\alpha_1 - r))z\partial_z V + \frac{1}{2}\sigma_2^2 y^2 \partial_{yy} V + \frac{1}{2}\sigma_1^2\beta^2 z^2 \partial_{zz} V + \beta\rho\sigma_1\sigma_2 yz\partial_{yz} V \\ &+ \sup_{0\leq\pi\leq 1} \big\{ \big[(\alpha_1 - r)x\partial_x V + \beta\sigma_1^2 xz\partial_{xz} V + \rho\sigma_1\sigma_2 xy\partial_{xy} V\big]\pi + \frac{1}{2}\sigma_1^2 x^2 \partial_{xx} V\pi^2\big\} \end{aligned} \tag{17}$$

where $\pi_t = \xi_t / X_t$ is the fraction of X_t invested in the liquid stock.

Next, we specify the boundary conditions. When $X_t = 0$, the manager cannot buy the illiquid stock due to the leverage constraint. Therefore, when $x = 0$ we must have:

$$\min\left\{-\partial_t V - \mathcal{L}V,\ \ \partial_y V - (1-\mu)\partial_x V\right\}|_{x=0} = 0 \tag{18}$$

Similarly, when $Y_t = 0$, the manager cannot sell the illiquid stock due to the short-selling constraint. Therefore, at the boundary $y = 0$, we must have that:

$$\min\left\{-\partial_t V - \mathcal{L}V,\ \ (1+\lambda)\partial_x V - \partial_y V\right\}|_{y=0} = 0 \tag{19}$$

The homogeneity of the CRRA preferences, and the linearity of (2), (3), (4) and (5) from section 2, imply that $V(t, ax, ay, az) = a^{1-\gamma}V(t, x, y, z)$ for any $a > 0$. Thus, by taking $a = \frac{1}{x+y}$, we obtain:

$$V\left(t, \frac{x}{x+y}, \frac{y}{x+y}, \frac{z}{x+y}\right) = \frac{1}{(x+y)^{1-\gamma}}V(t, x, y, z) \tag{20}$$

Thus, the solution to our problem can be characterized by a three-dimensional state variable (t, ζ, η), where $\zeta = \frac{y}{x+y}$ is the portfolio weight of the illiquid stock, and $\eta = \ln\frac{x+y}{z}$ is the fund's performance relative to the benchmark. We denote the left hand side of (20) by the following:

$$h(t, \zeta, \eta) = V(t, 1-\zeta, \zeta, e^{-\eta}) \tag{21}$$

and further define

$$\varphi(t, \zeta, \eta) = \frac{1}{1-\gamma}\log[(1-\gamma)h(t, \zeta, \eta)] \tag{22}$$

which leads to

$$V(t, x, y, z) = \frac{(x+y)^{1-\gamma}}{1-\gamma}e^{(1-\gamma)\varphi(t,\zeta,\eta)} = \frac{[(x+y)e^{\varphi(t,\zeta,\eta)}]^{1-\gamma}}{1-\gamma} \tag{23}$$

where $\varphi(t, \zeta, \eta)$ is the compounded interest rate that makes the manager indifferent between following her optimal investment policy or receiving this fixed rate on her initial wealth $x+y$, i.e., it represents the certainty equivalent rate of return.

In order to derive the equation that governs $\varphi(t, \zeta, \eta)$, we use the chain rule to calculate the partial derivatives of V through the partial derivatives of φ. We then obtain that $\varphi(t,\zeta,\eta)$ satisfies the following equation:

$$\min\{-\varphi_t - \mathcal{M}\varphi, \ \ \mathcal{M}_1\varphi, \ \ \mathcal{M}_2\varphi\} = 0 \tag{24}$$

in $\Sigma = \{(t,\zeta,\eta) : 0 \le t \le T, 0 \le \zeta \le 1, \eta \in \mathbb{R}\}$, where

$$\begin{aligned}
\mathcal{M}\varphi &= a_0 + a_1\varphi_\zeta + a_2\varphi_\eta + a_3(\varphi_{\zeta\zeta} + \gamma'\varphi_\zeta^2) + a_4(\varphi_{\eta\eta} + \gamma'\varphi_\eta^2) + a_5(\varphi_{\zeta\eta} + \gamma'\varphi_\zeta\varphi_\eta) \\
&\quad + \sup_{0\le\pi\le1}\Big\{ \left[b_0 + b_1\varphi_\zeta + b_2\varphi_\eta + b_3(\varphi_{\zeta\zeta} + \gamma'\varphi_\zeta^2) + b_4(\varphi_{\eta\eta} + \gamma'\varphi_\eta^2) + b_5(\varphi_{\zeta\eta} + \gamma'\varphi_\zeta\varphi_\eta)\right]\pi \\
&\quad + \left[c_0 + c_1\varphi_\zeta + c_2\varphi_\eta + c_3(\varphi_{\zeta\zeta} + \gamma'\varphi_\zeta^2) + c_4(\varphi_{\eta\eta} + \gamma'\varphi_\eta^2) + c_5(\varphi_{\zeta\eta} + \gamma'\varphi_\zeta\varphi_\eta)\right]\pi^2 \Big\}, \\
\mathcal{M}_1\varphi &= \mu + (1-\mu\zeta)\phi_\zeta + \mu\phi_\eta, \\
\mathcal{M}_2\varphi &= \lambda - (1+\lambda\zeta)\phi_\zeta + \lambda\phi_\eta,
\end{aligned}$$

where $\gamma' = 1-\gamma$, and the remaining coefficients are as follows:

$$\begin{aligned}
& a_0 = r + (\alpha_2 - r)\zeta - \frac{1}{2}\gamma\sigma_2^2\zeta^2, \quad a_1 = \zeta(1-\zeta)(\alpha_2 - r - \gamma\sigma_2^2\zeta), \\
& a_2 = -\beta(\alpha_1 - r) - \zeta(r - \alpha_2 + \sigma_2^2\gamma\zeta - \beta\rho\sigma_1\sigma_2\gamma) + \frac{1}{2}\left(\sigma_2^2\zeta^2 + \sigma_1^2\beta^2 - 2\beta\rho\sigma_1\sigma_2\zeta\right), \\
& a_3 = \frac{1}{2}\sigma_2^2\zeta^2(1-\zeta)^2, \quad a_4 = \frac{1}{2}\left(\sigma_2^2\zeta^2 + \sigma_1^2\beta^2 - 2\beta\rho\sigma_1\sigma_2\zeta\right), \quad a_5 = -\zeta(1-\zeta)(\beta\rho\sigma_1\sigma_2 - \sigma_2^2\zeta), \\
& b_0 = (\alpha_1 - r - \rho\sigma_1\sigma_2\gamma\zeta)(1-\zeta), \quad b_1 = -(\alpha_1 - r)\zeta(1-\zeta) + \rho\sigma_1\sigma_2\gamma\zeta(1-\zeta)(2\zeta - 1), \\
& b_2 = -(1-\zeta)(r - \alpha_1 - \gamma\beta\sigma_1^2 + 2\gamma\rho\sigma_1\sigma_2\zeta) + (1-\zeta)(\rho\sigma_1\sigma_2\zeta - \beta\sigma_1^2), \\
& b_3 = -\rho\sigma_1\sigma_2\zeta^2(1-\zeta)^2, \quad b_4 = (1-\zeta)(\rho\sigma_1\sigma_2\zeta - \beta\sigma_1^2), \quad b_5 = -\zeta(1-\zeta)\left[\kappa\sigma_1\sigma_2(2\zeta-1) - \beta\sigma_1^2\right], \\
& c_0 = -\frac{1}{2}\sigma_1^2\gamma(1-\zeta)^2, \quad c_1 = \sigma_1^2\gamma\zeta(1-\zeta)^2, \quad c_2 = -\sigma_1^2\gamma(1-\zeta)^2 + \frac{1}{2}\sigma_1^2(1-\zeta)^2, \\
& c_3 = \frac{1}{2}\sigma_1^2\zeta^2(1-\zeta)^2, \quad c_4 = \frac{1}{2}\sigma_1^2(1-\zeta)^2, \quad c_5 = -\sigma_1^2\zeta(1-\zeta)^2.
\end{aligned}$$

The terminal condition is given by

$$\varphi(T,\zeta,\eta) = \ln(1 + f(\eta)). \tag{25}$$

Given the solution to equation (24) with terminal condition (25), for any given time $t \in [0,T]$, the spacial solution domain $\Sigma_t = \{(\zeta, \eta) : 0 \leq \zeta \leq 1, \eta \in \mathbb{R}\}$ splits into three regions:

(i) sell region:

$$SR \equiv \{(\zeta, \eta) : \mathcal{M}_1 \varphi = 0\};$$

(ii) buy region:

$$BR \equiv \{(\zeta, \eta) : \mathcal{M}_2 \varphi = 0\};$$

(iii) no-trading region:

$$NTR \equiv \{(\zeta, \eta) : \varphi_t + \mathcal{M}\varphi = 0\}.$$

Numerical Procedure:

We briefly explain the numerical technique used to solve the variational inequality described above. We apply the standard penalty methods described in Dai and Zhong (2010). Instead of directly solving equation (24), we consider the following penalty approximation:

$$\varphi_t + \mathcal{M}\varphi + K(-\mathcal{M}_1\varphi)^+ + K(-\mathcal{M}_2\varphi)^+ = 0, \tag{26}$$

where K is a large penalty parameter. In the main algorithm, we apply an iterative method with error tolerance $tol > 0$, on a standard finite differences grid, with the following steps (assume the function value at time $t + \Delta t$ is known):

Step 1: Let $i = 0$, make an initial guess $\varphi^0(t, \zeta, \eta) = \varphi(t + \Delta t, \zeta, \eta)$.

Step 2: Find $\pi_i^* = \arg\max_{0 \leq \pi \leq 1} f(\pi, \varphi^i)$, where

$$\begin{aligned} f(\pi, \varphi) \quad = \quad & \Big\{ \left[b_0 + b_1\varphi_\zeta + b_2\varphi_\eta + b_3(\varphi_{\zeta\zeta} + \gamma'\varphi_\zeta^2) + b_4(\varphi_{\eta\eta} + \gamma'\varphi_\eta^2) + b_5(\varphi_{\zeta\eta} + \gamma'\varphi_\zeta\varphi_\eta)\right]\pi \\ & + \left[c_0 + c_1\varphi_\zeta + c_2\varphi_\eta + c_3(\varphi_{\zeta\zeta} + \gamma'\varphi_\zeta^2) + c_4(\varphi_{\eta\eta} + \gamma'\varphi_\eta^2) + c_5(\varphi_{\zeta\eta} + \gamma'\varphi_\zeta\varphi_\eta)\right]\pi^2 \Big\}. \end{aligned}$$

Step 3: Solve the discretized version of the following equation:[30]

$$\varphi_t^{i+1} + \mathcal{M}_0(\pi_i^*)\varphi^{i+1} + K(-\mathcal{M}_1\varphi^{i+1})^+ + K(-\mathcal{M}_2\varphi^{i+1})^+ = 0, \tag{27}$$

where the operator $\mathcal{M}_0(\pi)\varphi$ is

$$\mathcal{M}_0(\pi)\varphi = a_0 + a_1\varphi_\zeta + a_2\varphi_\eta + a_3(\varphi_{\zeta\zeta} + \gamma'\varphi_\zeta^2) + a_4(\varphi_{\eta\eta} + \gamma'\varphi_\eta^2) + a_5(\varphi_{\zeta\eta} + \gamma'\varphi_\zeta\varphi_\eta) + f(\pi, \varphi).$$

Step 4: If the following condition holds,

$$\frac{|\varphi^{i+1} - \varphi^i|}{\max\{1, |\varphi^i|\}} < \epsilon,$$

then we set $\varphi(t, \zeta, \eta) = \varphi^{i+1}(t, \zeta, \eta)$. Otherwise, we set $i = i + 1$ and we go back to Step 2.

Appendix B. Discussion on Equilibrium

In our model, we have assumed that fund managers take stock prices as given, and we compute liquidity premia as the extra return that they would require to be indifferent between trading the illiquid stock and trading its perfectly liquid counterpart. We derive our empirical predictions from comparative statics analyses.

However, it would be interesting to extend this setting to allow for multiple fund managers with heterogeneous incentives who can trade with one another and who can determine stock prices endogenously. It is beyond the scope of this paper to provide such a model. We only provide a brief discussion of its potential structure and the challenges that it would entail.

Assume a two-fund model in which the fund managers are endowed with different benchmarks to cater to two different investors with different preferences for liquid and illiquid assets. The more conservative benchmark would focus on the liquid asset, and the aggressive benchmark would focus on the illiquid asset.

[30]When dealing with the nonlinear terms, Newton's iterative method (smooth or nonsmooth) is used (cf. Forsyth and Vetzal (2002)).

The fund managers would be given compensation contracts with a year-end bonus component. This would give the fund managers the motive to trade more frequently. In fact, the incentive to deviate from the benchmark, and given the disparity in benchmarks, could generate trades in the same asset but in opposite directions, like in the model of Goncalves-Pinto, Sotes-Paladino, and Xu (2018). Specifically, the manager following a liquid benchmark would deviate by taking bets with the illiquid asset, while the manager following the illiquid benchmark would deviate by taking bets with the liquid asset. This would create trading opportunities.

It would be difficult to solve such a model, because of the complex interdependencies between the investment policies of the fund managers. However, we believe that the effect of convex incentives on the relation between trading costs and expected stock returns would be qualitatively similar to those reported in this paper. We leave this alternative framework for future research.

In a recent paper by Buss and Dumas (2019), they propose an algorithm to synchronize trades in a general-equilibrium setting with trading fees. They fully characterize the equilibrium and show that asset prices are not affected by the payment of the fees itself, but rather by the trade-off between smoothing consumption and smoothing holdings that the traders face. We believe that adding convex incentives to their model could potentially strengthen the effect of fee payments on asset prices.

Appendix C. Construction of Risk-Shifting Proxies

We follow Huang, Sialm, and Zhang (2011), who study the performance consequences of mutual funds varying the risk of their portfolios significantly over time. They propose a holdings-based measure to capture risk-shifting propensity. Specifically, they compare the risk of their current holdings, based on the fund's most recently disclosed positions, with the realized risk of the fund. They use rolling windows of 12 quarters, and several different measures of risk for the funds and their holdings. We describe these different measures of risk below. We use the ratio of current holdings risk to the fund's risk as a proxy for risk-shifting propensity, but the results would be qualitatively similar if we were to use the difference instead of the ratio.

All holdings volatility: This proxy uses all the portfolio holdings in the calculation of the risk ratio. The numerator is the standard deviation of the returns of all the portfolio holdings (including equity, bond, cash and others), over the prior 12 quarters, and the denominator is the standard deviation of the fund's realized returns over the same prior 12 quarters. If the ratio is larger than 1, then the fund is considered to be increasing the risk of its portfolio. The return of bonds and preferred stocks is considered to be the total return of the Barclay Capital Aggregate Bond Index, and for cash holdings and other assets, we use the Treasury bill rate as return.

Proportion of non-equity positions: Funds can shift portfolio risk by switching between equity and non-equity holdings, where equity is assumed to be riskier. We aggregate the portfolio proportions invested in cash, bonds, and other non-equity positions, over the prior 12 quarters, and divide this aggregate by the most recently disclosed non-equity aggregate portfolio proportion. If this ratio is smaller than 1, it means that the fund is decreasing the proportion of non-equity holdings in its portfolio, which corresponds to taking less risk.

Equity holdings volatility: In this proxy, we consider only the riskiness of the equity positions and ignore the non-equity positions. We compare the riskiness of the equity positions disclosed in the most recent disclosure quarter, with the riskiness of a hypothetical portfolio that maintains the historically disclosed positions in equity holdings. If the ratio is larger than 1, the fund is considered to be increasing the risk of its portfolio.

CAPM beta: It could be the case that fund managers change only the systematic risk of their portfolios, by switching between low beta equities and high beta equities. We estimate the market beta of every equity holding using the CAPM over the 12-quarter lookback window, and compare the betas of the equity holdings from the most recently disclosed portfolio, with the CAPM betas of the equity positions from the historically disclosed portfolios.

CAPM idiosyncratic volatility: This is the standard deviation of the residuals from the CAPM model used in the previous measure, using the same 12-quarter lookback window. We take the CAPM idiosyncratic risk from the most recently disclosed portfolio and divide it by its counterpart using the historically disclosed portfolios. This ratio is larger than 1 if the fund is increasing its CAPM idiosyncratic risk.

Carhart idiosyncratic volatility: For this proxy, we compute idiosyncratic risk using the standard deviation of the residuals from the Carhart four-factor model. The treatment of this measure is otherwise similar to that used for its CAPM counterpart.

Tracking error (value-weighted): The tracking error volatility is the standard deviation of the difference between the fund (or holdings) returns and the benchmark return. For this proxy, we use the value-weighted total market return from CRSP as the benchmark return. We take the tracking error of the most recently disclosed holdings and divide it by the tracking error of the fund's realized returns. If this ratio is larger than 1, then the fund is increasing its tracking error volatility.

Tracking error (equal-weighted): For this proxy, tracking error volatility is computed using the equal-weighted total market return from CRSP as the benchmark return. The treatment of this measure is otherwise similar to that for its value-weighted counterpart.

Table 1: Trading Characteristics and the Optimal Investment Policy

This table provides the results on some statistics of the optimal trading policy, including: the average number of trades executed over the investment horizon (Number of Trades), the total volume of trading over the investment horizon as a fraction of the fund's initial AUM (Total Volume (%)), the present value of the transaction costs paid as a fraction of the fund's initial AUM (PVTC(%)), and the expected duration from purchase to sale (Time from Buy to Sell). We obtained these results from 10,000 Monte Carlo simulations of the optimal investment policy in our model. In the simulations, we assume at most two trades per business day. When calculating the expected time duration from purchase to sale, we restrict our attention to the sample paths along which there is at least one purchase and one sale (note that in Panel B and C, no such paths are found in the case with no convex incentives). We report the results for both the cases with or without bonuses. The parameter values used to generate these results are as follows: the managerial risk aversion coefficient is $\gamma = 5$; the fund manager's investment horizon is $T = 1$ year; the risk-free rate is $r = 0.04$; the expected value and volatility of the returns on the liquid benchmark stock are $\alpha_1 = 0.09$ and $\sigma_1 = 0.14$; the expected value and volatility of the returns on the illiquid non-benchmark stock are $\alpha_2 = 0.19$ and $\sigma_2 = 0.24$; the return correlation between the two stocks is $\rho = 0.53$; the benchmark is assumed to solely consist of the liquid stock, i.e., $\beta = 1$. For the case with bonuses, the parameters in the bonus-performance function are matched to the empirical estimates of Lee, Trzcinka, and Venkatesan (2019) and Ma, Tang, and Gomez (2019), as follows: $\theta_L = 0.01$, $\theta_H = 0.15$, $f_L = 0$, and $f_H = 1.5$.

	Number of Trades	Total Volume(%)	PVTC (%)	Time from Buy to Sell
Panel A: $\lambda = \mu = 0.005$				
No Bonus Case	3.533	0.752	0.004	0.368
Bonus Case	29.482	63.060	0.309	0.161
Panel B: $\lambda = \mu = 0.01$				
No Bonus Case	0.360	0.083	0.001	N.A.
Bonus Case	17.973	49.089	0.481	0.181
Panel C: $\lambda = \mu = 0.02$				
No Bonus Case	0.001	0.000	0.000	N.A.
Bonus Case	6.647	25.143	0.495	0.392

Table 2: Optimal Policy and Liquidity Premia

This table provides information on the optimal trading policy at the initial time $t = 0$, and on the liquidity premia commanded by the fund manager, for multiple values of the transaction cost rate. Panel A reports the results for the case without convex incentives, and Panel B represents the case with bonuses. In Panel A, $S(0)$ and $B(0)$ are the levels of the sell boundary and of the buy boundary, which are independent of the performance of the fund relative to the benchmark due to the absence of convex incentives. In Panel B, $S^*(0,\eta)$ and $B^*(0,\eta)$ ($S_*(0,\eta)$ and $B_*(0,\eta)$) are the max (min) levels of the sell and the buy boundaries across different values of relative performance (η). In Panel A, δ_c is the liquidity premium commanded by the fund manager, i.e., the maximum level of expected return on the illiquid non-benchmark stock that the fund manager is willing to forego in exchange for zero transaction costs. In Panel B, it is denoted as δ. The variable δ_c^0 (δ^0) is the liquidity premium exclusively due to the suboptimal risk exposure due to the presence of transaction costs. The parameter values used to generate these results are as follows: the managerial risk aversion coefficient is $\gamma = 5$; the fund manager's investment horizon is $T = 1$ year; the risk-free rate is $r = 0.04$; the expected value and volatility of the returns on the liquid benchmark stock are $\alpha_1 = 0.09$ and $\sigma_1 = 0.14$; the expected value and volatility of the returns on the illiquid non-benchmark stock are $\alpha_2 = 0.19$ and $\sigma_2 = 0.24$; the return correlation of the two stocks is $\rho = 0.53$; the benchmark is assumed to solely consist of the liquid stock, i.e., $\beta = 1$. For the case with bonuses (Panel B), the parameters values for the bonus-performance function are matched to the empirical estimates in Lee, Trzcinka, and Venkatesan (2019) and Ma, Tang, and Gomez (2019), as follows: $\theta_L = 0.01$, $\theta_H = 0.15$, $f_L = 0$, and $f_H = 1.5$.

Panel A: No Bonus Case

$\lambda = \mu =$	0.005	0.01	0.02	0.03	0.04	0.05
$S(0)$	0.56	0.58	0.62	0.65	0.69	0.73
$B(0)$	0.45	0.43	0.38	0.34	0.29	0.25
δ_c (%)	0.031	0.034	0.034	0.034	0.034	0.034
δ_c^0 (%)	0.022	0.032	0.034	0.034	0.034	0.034
$\delta_c/(\lambda+\mu)$	0.031	0.017	0.009	0.006	0.004	0.003
$\delta_c^0/\delta_c \times 100$	73.02	94.68	100.00	100.00	100.00	100.00

Panel B: Bonus Case

$\lambda = \mu =$	0.005	0.01	0.02	0.03	0.04	0.05
$S^*(0,\eta)$	0.72	0.74	0.77	0.8	0.82	0.84
$S_*(0,\eta)$	0.05	0.05	0.06	0.06	0.07	0.08
$B^*(0,\eta)$	0.57	0.55	0.51	0.44	0.35	0.28
$B_*(0,\eta)$	0.01	0.00	0.00	0.00	0.00	0.00
δ (%)	1.507	2.349	3.579	4.204	4.534	4.684
δ^0 (%)	0.724	1.193	2.551	3.571	4.284	4.613
$\delta/(\lambda+\mu)$	1.507	1.175	0.895	0.701	0.567	0.468
$\delta^0/\delta \times 100$	48.00	50.77	71.29	84.93	94.50	98.49
δ/δ_c	49.22	69.03	104.28	122.50	132.10	136.49

Table 3: Comparative Statics

This table provides information on the optimal trading policy at the initial time $t = 0$, and on the liquidity premia commanded by the fund manager, for multiple values of the model parameters. Panel A reports the results for the case without bonuses, and Panel B represents the case with bonuses. In Panel A, $S(0)$ and $B(0)$ are the levels of the sell and the buy boundaries, which are independent of the performance of the fund relative to the benchmark due to the absence of convex incentives. In Panel B, $S(0,0)$ and $B(0,0)$ are the levels of the sell and buy boundaries at the initial time $t = 0$, with the value of relative performance set to zero ($\eta = 0$). In Panel A, δ_c is the liquidity premium commanded by the fund manager, i.e., the maximum level of expected return on the illiquid non-benchmark stock that the fund manager is willing to forego in exchange for zero transaction costs. In Panel B it is denoted as δ. The parameter values used to generate the baseline results are as follows: the managerial risk aversion coefficient is $\gamma = 5$; the fund manager's investment horizon is $T = 1$ year; the risk-free rate is $r = 0.04$; the expected value and volatility of the returns on the liquid benchmark stock are $\alpha_1 = 0.09$ and $\sigma_1 = 0.14$; the expected value and volatility of the returns on the illiquid non-benchmark stock are $\alpha_2 = 0.19$ and $\sigma_2 = 0.24$; the return correlation of the two stocks is $\rho = 0.53$; the benchmark is assumed to solely consist of the liquid stock, i.e., $\beta = 1$. For the case with bonuses (Panel B), the parameters values for the bonus-performance function are matched to the empirical estimates in Lee, Trzcinka, and Venkatesan (2019) and Ma, Tang, and Gomez (2019), as follows: $\theta_L = 0.01$, $\theta_H = 0.15$, $f_L = 0$, and $f_H = 1.5$. In each panel, the transaction costs rates are $\lambda = \mu = 0.01$.

	Panel A: No Bonus Case			Panel B: Bonus Case			
	$S(0)$	$B(0)$	$\frac{\delta_c}{\lambda+\mu}$	$S(0,0)$	$B(0,0)$	$\frac{\delta}{\lambda+\mu}$	$\frac{\delta}{\delta_c}$
Base case	0.58	0.43	0.017	0.69	0.21	1.175	69.03
$\gamma \times 1.1$	0.53	0.39	0.020	0.60	0.17	1.221	59.99
$\gamma \times 0.9$	0.64	0.48	0.013	0.80	0.26	1.111	84.13
$\rho \times 1.1$	0.59	0.44	0.020	0.69	0.21	1.138	55.62
$\rho \times 0.9$	0.56	0.42	0.018	0.69	0.20	1.207	68.59
$\alpha_2 \times 1.1$	0.65	0.52	0.018	0.74	0.25	1.129	61.39
$\alpha_2 \times 0.9$	0.49	0.34	0.018	0.63	0.15	1.238	68.58
$\sigma_2 \times 1.1$	0.47	0.33	0.027	0.59	0.16	1.314	48.70
$\sigma_2 \times 0.9$	0.71	0.57	0.010	0.82	0.27	1.040	102.68
$\beta \times 0.9$	0.58	0.43	0.017	0.72	0.23	1.172	68.90
$\beta \times 0.8$	0.58	0.43	0.017	0.75	0.25	1.168	68.65
$\theta_H + 1\%$	0.58	0.43	0.017	0.69	0.21	1.144	67.25
$\theta_H - 1\%$	0.58	0.43	0.017	0.69	0.20	1.207	70.95
$f_H = 2.5$	0.58	0.43	0.017	0.68	0.16	1.401	82.34
$f_H = 0.5$	0.58	0.43	0.017	0.68	0.31	0.713	41.90

Table 4: Portfolio Statistics

This table reports summary statistics for 18 out of 50 portfolios of stocks used in our analysis. Assignment of a stock to a particular portfolio in a given test year depends on three criteria: (1) the average RSI in the previous year (two groups), (2) the Gibbs Beta in the previous year (five groups), and (3) the Gibbs c in the previous year (five groups). The Gibbs Beta and c measures are estimated following Hasbrouck (2009). In Panel A, the RSI measure is computed using APR, which is the average percentile rank across the all the stock-level proxies. In December of the prior year, we assign percentile ranks to each stock-level proxy for the whole cross-section of stocks. Next, we compute the average percentile rank across the proxies for a given stock. In Panel B, the RSI measure is computed as FPC, which is the first principal component across all stock-level proxies from Huang, Sialm, and Zhang (2011). We report both groups of RSI, but only the quintiles 1 (Low), 3 (Mid), and 5 (high) of Beta and c. The table reports time-series averages of equal-weighted portfolio means. The sample covers the period 2004-2017.

Panel A: Average Percentile Ranks

APR	Beta Rank	c Rank	Ret	Beta	c	LRMC	# Stocks	Turnover
Below-Median	Low	Low	0.0058	0.3697	0.0036	-0.8525	34	0.5511
		Mid	0.0062	0.3345	0.0089	-1.2858	35	0.4900
		High	0.0125	0.3062	0.0221	-2.1328	34	0.5862
	Mid	Low	0.0055	0.9186	0.0021	1.1626	35	1.1228
		Mid	0.0068	0.8564	0.0049	-0.3693	35	1.1606
		High	0.0091	0.6861	0.0145	-1.6287	33	1.1676
	High	Low	0.0070	1.4209	0.0026	1.3462	36	1.7232
		Mid	0.0103	1.4890	0.0047	0.4184	35	2.0072
		High	0.0025	1.3576	0.0095	-0.3038	34	2.1495
Above-Median	Low	Low	0.0065	0.7964	0.0016	1.9815	35	1.5712
		Mid	0.0069	0.7814	0.0035	0.3832	35	1.5073
		High	0.0137	0.5298	0.0108	-1.1912	35	0.8678
	Mid	Low	0.0051	1.1453	0.0016	2.4809	36	1.9598
		Mid	0.0052	1.1928	0.0028	1.4811	36	2.2065
		High	0.0103	1.2001	0.0054	0.4984	36	2.0397
	High	Low	0.0076	1.4966	0.0023	1.8912	37	2.3715
		Mid	0.0096	1.6078	0.0039	1.1861	36	3.0415
		High	0.0081	1.5870	0.0071	0.6223	36	2.5854

Panel B: First Principal Component

FPC	Beta Rank	c Rank	Ret	Beta	c	LRMC	# Stocks	Turnover
Below-Median	Low	Low	0.0053	0.3640	0.0036	-0.8550	34	0.5356
		Mid	0.0065	0.3395	0.0089	-1.2590	35	0.4976
		High	0.0134	0.3015	0.0220	-2.1073	34	0.5855
	Mid	Low	0.0065	0.9153	0.0021	1.1499	35	1.1296
		Mid	0.0085	0.8528	0.0049	-0.3800	35	1.1921
		High	0.0096	0.6897	0.0143	-1.6235	33	1.1976
	High	Low	0.0075	1.4312	0.0026	1.3102	36	1.6894
		Mid	0.0091	1.4778	0.0047	0.4043	36	1.9745
		High	0.0033	1.3651	0.0095	-0.2974	34	2.1099
Above-Median	Low	Low	0.0067	0.8018	0.0016	1.9925	35	1.5780
		Mid	0.0112	0.7831	0.0036	0.3081	35	1.5194
		High	0.0136	0.5075	0.0115	-1.2875	35	0.8728
	Mid	Low	0.0053	1.1430	0.0016	2.4804	36	1.9438
		Mid	0.0061	1.1863	0.0028	1.5082	35	2.2534
		High	0.0114	1.2023	0.0055	0.4774	36	2.0752
	High	Low	0.0068	1.4926	0.0023	1.8957	37	2.4093
		Mid	0.0088	1.6089	0.0039	1.1447	36	2.9779
		High	0.0076	1.5991	0.0071	0.6194	36	2.5898

Table 5: Convex Incentives and Liquidity Premia

This table reports the GMM results following the methodology in Hasbrouck (2009). In this analysis, we include only stocks from NASDAQ. We form 50 portfolios by sequentially sorting stocks based on the average value of RSI in the previous year (two groups), the beta estimates from the prior three years (five groups), and effective trading cost (five groups). We compute the RSI measure using two methods. First, we use the average percentile rank across the stock-level proxies, i.e., APR. The second method for computing RSI uses principal component analysis. In December of the prior year, we compute the principal components across the stock-level proxies, and use the first principal component, i.e., FPC. We use the 8 risk-shifting proxies from Huang, Sialm, and Zhang (2011). For portfolio formation, both beta and effective cost estimates are the Gibbs estimates of the basic market-factor model in Hasbrouck (2009). The dependent variable is the (equally-weighted) monthly stock return for each portfolio in the year after formation. In Panels A and B we control for *MKT Beta*, *SMB Beta*, and *HML Beta*, which are the unconditional betas obtained from a three-factor model of Fama and French (1993), estimated over the entire sample period for each portfolio. In Panel C we include additional controls, such as *UMD Beta* from the four-factor model of Carhart (1997), and *RMW Beta* and *CMA Beta* from the five-factor model of Fama and French (2015). *Turnover* is the portfolio average ratio of stock trading volume to the number of shares outstanding, *LRMC* is the log relative market capitalization (i.e., the average median-adjusted market capitalization for the stocks in each portfolio), and *Ownership* is the number of shares held by active funds divided by the shares outstanding. The effective trading cost measure is denoted as c. To assess the impact of risk-shifting on the relation between trading costs and future stock returns, the regressions include the indicator $DummyRSI$ and its interaction with the effective trading cost. The indicator $DummyRSI$ is equal to one for portfolios with above-median values of RSI, and equals zero otherwise. We compute t-statistics (reported in parentheses) using GMM standard errors that correct for estimation error in the unconditional betas and for heteroskedasticity, and $*$, $**$, and $***$ represent significance at the 10%, 5%, and 1% levels, respectively.

	Panel A: Average Percentile Ranks (APR)					Panel B: First Principal Component (FPC)				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
c	1.8390***	1.8311***	1.7390***	1.0863***	1.0767***	1.8407***	1.8529***	1.9475***	1.1073***	1.2625***
	(8.75)	(8.65)	(7.89)	(4.39)	(4.24)	(8.76)	(8.58)	(8.90)	(4.46)	(5.02)
DummyRSI			-0.0113***		-0.0085***			-0.0314***		-0.0274***
			(-4.58)		(-3.34)			(-6.21)		(-5.35)
c × DummyRSI			1.4334***		1.0740***			2.2653***		1.6408***
			(3.87)		(2.96)			(4.03)		(2.95)
Turnover				-0.1375***	-0.1231***				-0.1437***	-0.1260***
				(-5.47)	(-4.86)				(-5.74)	(-4.99)
c × Turnover				14.5452***	13.3996***				14.5758***	13.7773***
				(4.68)	(4.37)				(4.70)	(4.47)
LRMC	0.0039***	0.0035**	0.0039***	0.0041***	0.0042***	0.0036***	0.0038***	0.0031**	0.0046***	0.0034**
	(4.53)	(2.51)	(2.64)	(2.89)	(2.90)	(3.96)	(3.02)	(2.23)	(3.64)	(2.46)
Ownership		0.0052	0.0226*	0.0393***	0.0488***		-0.0033	0.0975***	0.0279**	0.1193***
		(0.37)	(1.45)	(2.58)	(3.03)		(-0.29)	(4.45)	(2.29)	(5.43)
MKT Beta	0.0044	0.0040	0.0060	0.0068	0.0078	0.0057	0.0062	0.0051	0.0114*	0.0103
	(0.70)	(0.63)	(0.88)	(1.07)	(1.14)	(0.79)	(0.84)	(0.68)	(1.51)	(1.36)
SMB Beta	-0.0054	-0.0053	-0.0087*	-0.0097*	-0.0122**	-0.0059	-0.0060	-0.0080*	-0.0104*	-0.0128**
	(-1.13)	(-1.11)	(-1.76)	(-1.91)	(-2.29)	(-1.18)	(-1.21)	(-1.60)	(-1.89)	(-2.33)
HML Beta	0.0029	0.0027	0.0034	-0.0026	-0.0017	0.0035	0.0039	0.0065	0.0006	0.0034
	(0.47)	(0.43)	(0.52)	(-0.41)	(-0.26)	(0.55)	(0.60)	(0.98)	(0.09)	(0.52)
Intercept	-0.0011	-0.0017	-0.0014	0.0072	0.0069	-0.0020	-0.0019	-0.0090	0.0058	-0.0023
	(-0.17)	(-0.26)	(-0.21)	(1.10)	(0.97)	(-0.30)	(-0.29)	(-1.29)	(0.87)	(-0.33)
Observations	8,400	8,400	8,400	8,400	8,400	8,400	8,400	8,400	8,400	8,400

Table 5: Continued

Panel C: Alternative Factor Models

	APR				FPC			
	FF4		FF5		FF4		FF5	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
c	1.7531***	1.0626***	1.8857***	1.2148***	1.9622***	1.2584***	2.0864***	1.4053***
	(7.96)	(4.18)	(8.23)	(4.69)	(8.97)	(5.00)	(9.03)	(5.38)
DummyRSI	-0.0093***	-0.0057**	-0.0137***	-0.0111***	-0.0302***	-0.0251***	-0.0304***	-0.0262***
	(-3.25)	(-1.98)	(-4.73)	(-3.79)	(-5.64)	(-4.63)	(-5.51)	(-4.70)
c × DummyRSI	1.3446***	0.9391**	1.6172***	1.2714***	2.1648***	1.4187**	2.2800***	1.6530***
	(3.53)	(2.53)	(4.27)	(3.44)	(3.63)	(2.40)	(3.90)	(2.86)
Turnover		-0.1262***		-0.1254***		-0.1294***		-0.1260***
		(-4.98)		(-4.94)		(-5.12)		(-5.00)
c × Turnover		14.1249***		13.6743***		14.5010***		13.6798***
		(4.59)		(4.45)		(4.68)		(4.46)
LRMC	0.0042***	0.0048***	0.0022*	0.0025*	0.0032**	0.0036***	0.0013	0.0016
	(2.82)	(3.17)	(1.45)	(1.65)	(2.30)	(2.59)	(0.86)	(1.06)
Ownership	0.0196	0.0459***	0.0295*	0.0566***	0.0962***	0.1181***	0.0992***	0.1209***
	(1.25)	(2.83)	(1.88)	(3.46)	(4.39)	(5.37)	(4.52)	(5.50)
MKT Beta	-0.0001	-0.0009	0.0067	0.0089	0.0025	0.0053	0.0047	0.0098
	(-0.01)	(-0.10)	(0.92)	(1.22)	(0.29)	(0.62)	(0.62)	(1.26)
SMB Beta	-0.0093*	-0.0137**	-0.0013	-0.0052	-0.0085*	-0.0143**	-0.0011	-0.0063
	(-1.88)	(-2.54)	(-0.22)	(-0.85)	(-1.67)	(-2.55)	(-0.20)	(-1.01)
HML Beta	0.0034	-0.0020	-0.0107	-0.0173**	0.0073	0.0046	-0.0036	-0.0068
	(0.53)	(-0.31)	(-1.31)	(-2.08)	(1.09)	(0.70)	(-0.43)	(-0.82)
UMD Beta	-0.0162*	-0.0203*			-0.0120	-0.0203*		
	(-1.54)	(-1.93)			(-1.08)	(-1.82)		
RMW Beta			0.0167***	0.0176***			0.0152***	0.0152***
			(3.60)	(3.76)			(2.90)	(2.90)
CMA Beta			-0.0082*	-0.0110**			-0.0083*	-0.0110**
			(-1.54)	(-2.05)			(-1.51)	(-2.01)
Intercept	0.0015	0.0115*	-0.0044	0.0038	-0.0083	-0.0004	-0.0116*	-0.0048
	(0.19)	(1.46)	(-0.60)	(0.51)	(-1.15)	(-0.06)	(-1.61)	(-0.67)
Observations	8,400	8,400	8,400	8,400	8,400	8,400	8,400	8,400

Table 6: Risk-Shifting and Turnover

This table reports the results of a regression of portfolio turnover on lagged trading costs, lagged risk-shifting incentives, and their interaction. We use the methodology in Hasbrouck (2009). The dependent variable is the monthly turnover (i.e., the ratio of trading volume to the number of shares outstanding) of each portfolio in the year after formation (equal-weighted across the stocks in each portfolio). To assess the impact of the risk-shifting on the relation between trading costs and future turnover, the regressions include the indicator *DummyRSI* and its interaction with the effective trading cost. We control for several stock characteristics that are important determinants of stock turnover, following Section 4.1 of Lo and Wang (2000). Specifically, we control for (i) the natural log of a stock's market capitalization, averaged across all stocks in a portfolio (*Ln(Size)*), (ii) the intercept coefficient from the time-series regression of a stock's return on the value-weighted market return, averaged across all stocks in a portfolio (*Alpha*), (iii) the slope coefficient from the time-series regression of a stock's return on the value-weighted market return, averaged across all stocks in a portfolio (*Beta*), (iv) the residual standard deviation of the time-series regression of a stock's return on the value-weighted market return, averaged across all stocks in a portfolio (*IdioVol*), and (v) the average dividend yield of the stocks in each portfolio (*DivYield*). In some specifications, we also control for *Ownership*, which is the number of shares held by active funds divided by the shares outstanding. t-statistics are reported in parenthesis and are computed using GMM standard errors that correct for estimation error in the unconditional betas and for heteroskedasticity, and $*$, $**$, and $***$ represent significance at the 10%, 5%, and 1% levels, respectively.

	APR		FPC	
	(1)	(2)	(3)	(4)
c	0.5981*	0.1073	0.1226	-0.1305
	(1.81)	(0.34)	(0.27)	(-0.29)
DummyRSI	0.0489***	0.0350***	0.0714***	0.0417***
	(23.05)	(16.59)	(26.22)	(11.14)
c × DummyRSI	-2.5418***	-2.6847***	-4.0317***	-3.1371***
	(-7.82)	(-8.50)	(-10.19)	(-7.94)
Ln(Size)	0.0253***	0.0093***	0.0141***	0.0070***
	(26.61)	(7.30)	(11.66)	(5.04)
Alpha	0.4149	0.1572	0.7728	0.5143
	(1.32)	(0.49)	(1.34)	(0.88)
Beta	0.0652***	0.0600***	0.0841***	0.0771***
	(13.13)	(11.90)	(11.06)	(10.02)
IdioVol	0.3327***	0.3309***	0.1505***	0.1808***
	(9.34)	(9.17)	(2.94)	(3.52)
DivYield	-0.4351**	-0.5729***	-0.5455***	-0.6428***
	(-2.49)	(-3.45)	(-2.92)	(-3.49)
Ownership		0.2302***		0.1793***
		(17.17)		(10.23)
Intercept	-0.3090***	-0.1389***	-0.1658***	-0.0934***
	(-22.28)	(-8.62)	(-9.61)	(-5.06)
Observations	8,400	8,400	8,400	8,400

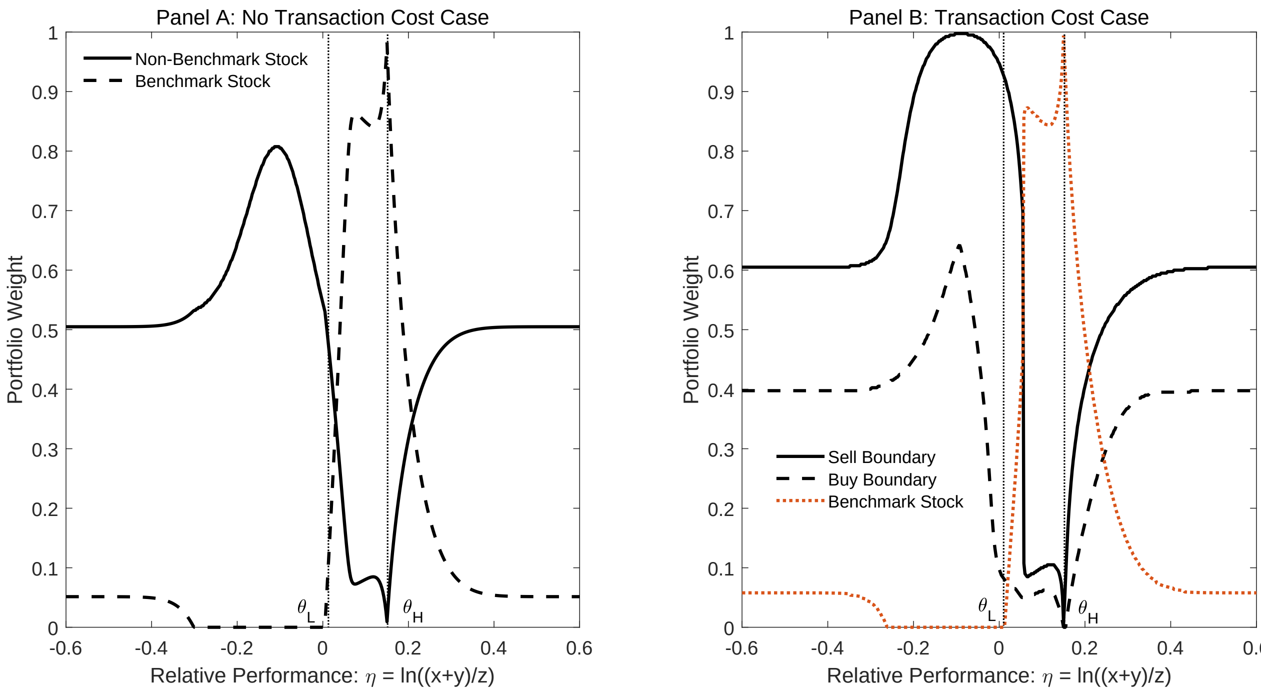

Figure 1: Optimal Stock Allocations

This figure shows the fund's optimal allocation to the benchmark stock and the non-benchmark stock, as a function of the fund's performance relative to the benchmark, when the bonuses have a collar specification. This is a snapshot of the policy at mid-year ($t = 0.5$). Panel A shows the case without transaction costs ($\lambda = \mu = 0$), and Panel B shows the case with a transaction cost rate of 1% ($\lambda = \mu = 1\%$). In Panel A, the fund trades continuously to maintain the optimal exposures on the benchmark stock (dashed line) and the non-benchmark stock (solid line). In Panel B, the fund only trades the non-benchmark stock when the allocation on this stock is either above the sell boundary (solid line) or below the buy boundary (dashed line). The dotted line in Panel B represents the average allocation on the benchmark stock. The parameter values used to generate these results are as follows: the managerial risk aversion coefficient is $\gamma = 5$; the fund manager's investment horizon is $T = 1$ year; the risk-free rate is $r = 0.04$; the expected value and volatility of the returns on the liquid benchmark stock is $\alpha_1 = 0.09$ and $\sigma_1 = 0.14$; the expected value and volatility of the returns on the illiquid non-benchmark stock is $\alpha_2 = 0.19$ and $\sigma_2 = 0.24$; the return correlation of the two stocks is $\rho = 0.53$; the benchmark is assumed to be fully invested in the liquid stock, i.e., $\beta = 1$; the parameters in the bonus-performance function are matched to the empirical estimates of Lee, Trzcinka, and Venkatesan (2019) and Ma, Tang, and Gomez (2019), as follows: $\theta_L = 0.01$, $\theta_H = 0.15$, $f_L = 0$, and $f_H = 1.50$.

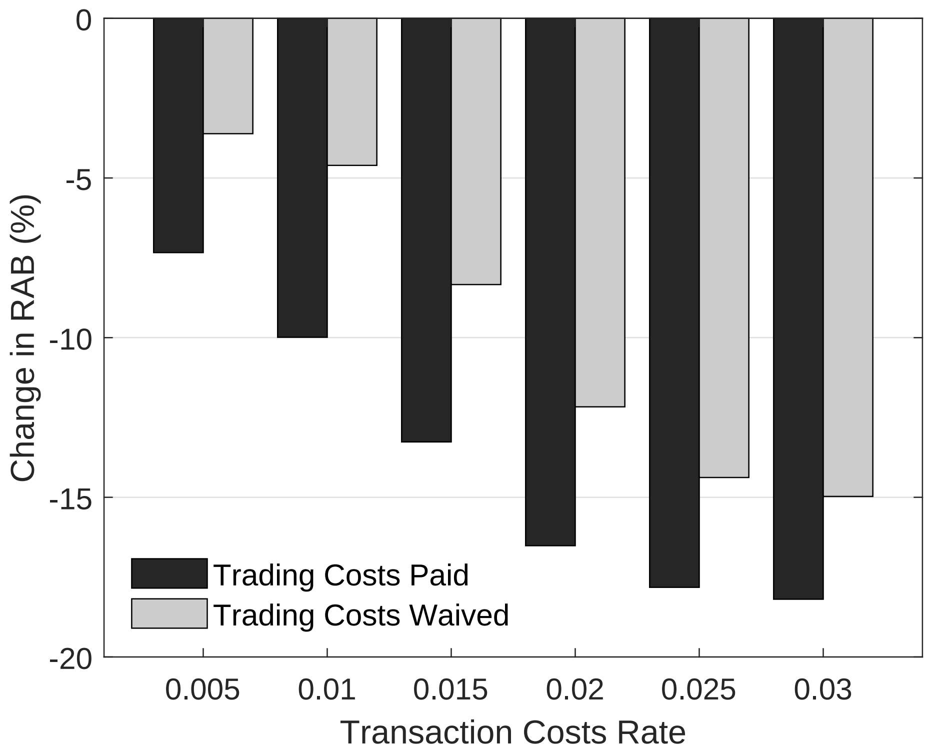

Figure 2: Trading Costs and Reduction in Risk-Adjusted Bonuses

This figure shows the decrease in the manager's risk-adjusted bonuses (RAB), which is defined as the minimum amount of additional AUM that the manager requires for waiving her bonuses, due to the presence of transaction costs on the non-benchmark stock. The black-coloured bars represent the total decrease in RAB, and the grey-coloured bars represent the decrease in RAB exclusively due to the suboptimal risk exposure caused by the presence of transaction costs (i.e., the trading costs are waived). The remaining parameter values used to generate these results are as follows: the managerial risk aversion coefficient is $\gamma = 5$; the fund manager's investment horizon is $T = 1$ year; the risk-free rate is $r = 0.04$; the expected return and the volatility of returns for the liquid benchmark stock are $\alpha_1 = 0.09$ and $\sigma_1 = 0.14$, respectively; the expected return and the volatility of returns for the illiquid non-benchmark stock are $\alpha_2 = 0.19$ and $\sigma_2 = 0.24$, respectively; the return correlation of the two stocks is $\rho = 0.53$; the benchmark is assumed to be fully invested in the liquid stock, i.e., $\beta = 1$; the parameters of the bonus function are as follows: $\theta_L = 0.01$, $\theta_H = 0.15$, $f_L = 0$, and $f_H = 1.5$.

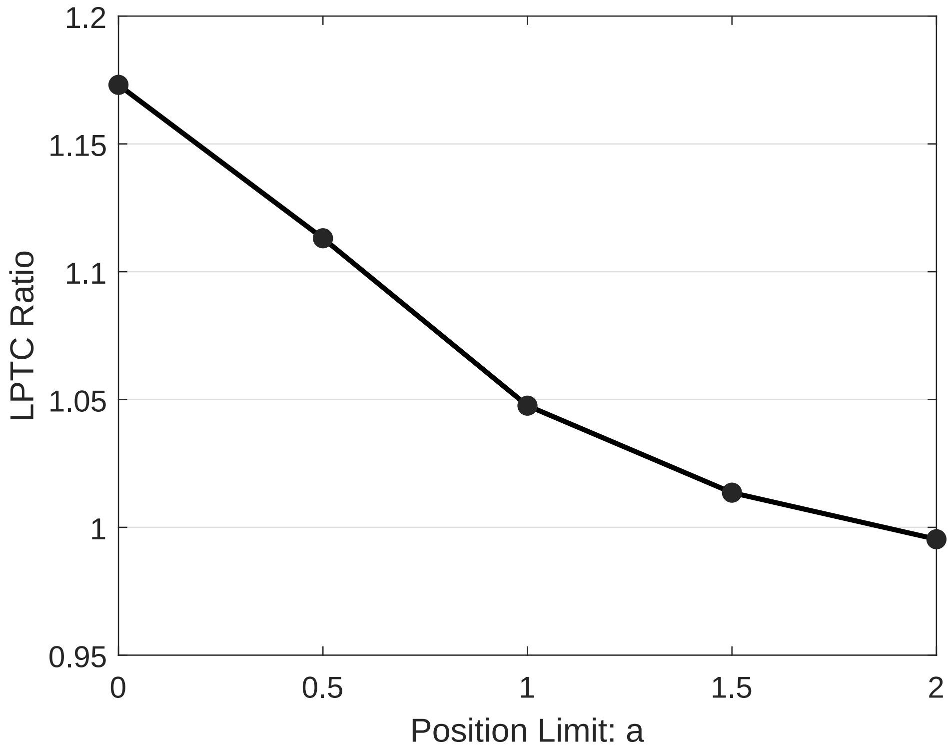

Figure 3: Relaxing the Position Limits on the Liquid Benchmark Stock

This figure shows the liquidity premium to transaction cost (LPTC) ratio for different position limits on the benchmark stock. We fix the position limit on the illiquid non-benchmark stock to be within the interval $[0, 1]$, but we relax the position limit on the liquid benchmark stock to fall in the interval $[-a, 1 + a]$, where a ranges from 0 to 2 (i.e., the x-axis in the figure). The remaining parameter values used to generate these results are as follows: the managerial risk aversion coefficient is $\gamma = 5$; the fund manager's investment horizon is $T = 1$ year; the risk-free rate is $r = 0.04$; the expected value and volatility of the returns on the liquid benchmark stock is $\alpha_1 = 0.09$ and $\sigma_1 = 0.14$; the expected return and the volatility of the returns of the illiquid non-benchmark stock are $\alpha_2 = 0.19$ and $\sigma_2 = 0.24$, respectively; the return correlation of the two stocks is $\rho = 0.53$; the benchmark is assumed to be fully invested in equity, i.e., $\beta = 1$; the parameters of the bonus-performance function are matched to the empirical estimates of Lee, Trzcinka, and Venkatesan (2019) and Ma, Tang, and Gomez (2019), as follows: $\theta_L = 0.01$, $\theta_H = 0.15$, $f_L = 0$, and $f_H = 1.5$.

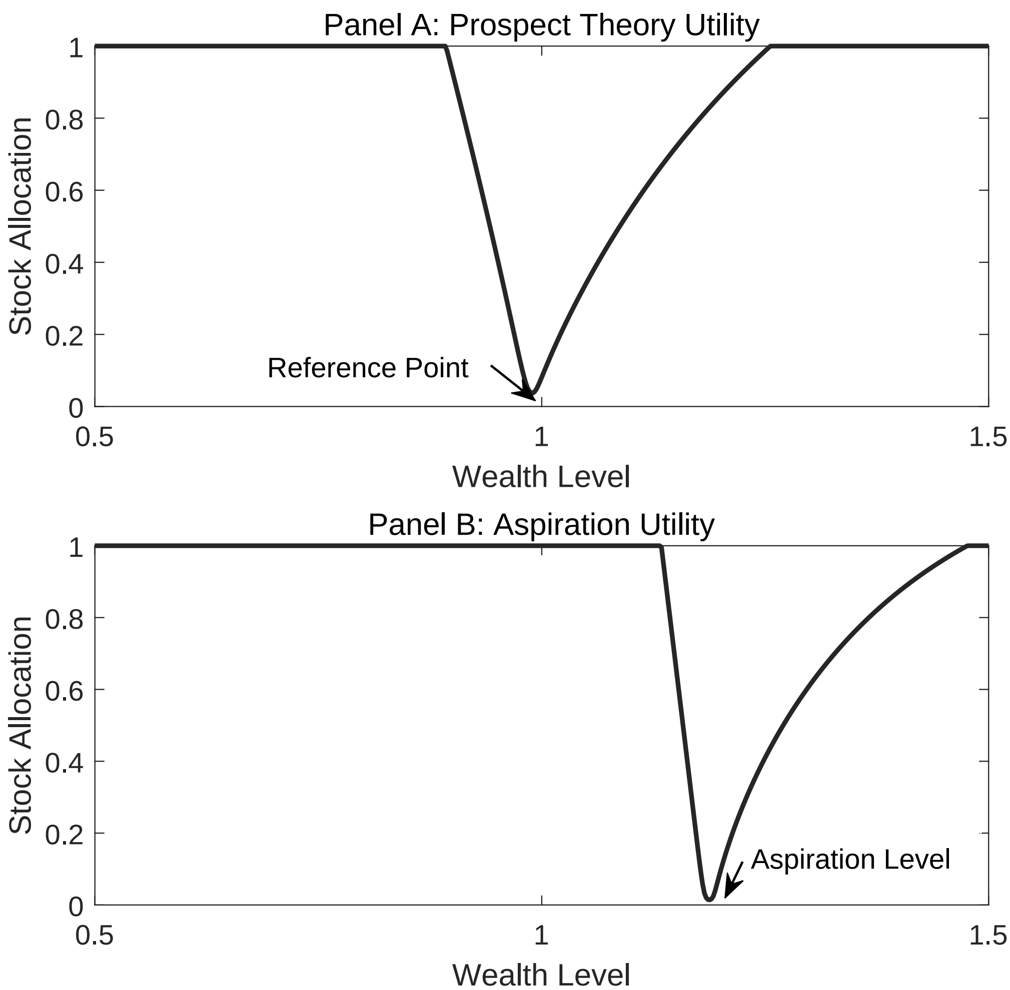

Figure 4: Optimal Trading Strategy with Non-Concave Utility Functions

This figure shows the optimal allocation in the stock at time $t = 0.5$, as a function of the investor's wealth level, when using the prospect theory utility function (Panel A) or the aspiration utility (Panel B). The parameter values used to generate these results are as follows: the investor's investment horizon is $T = 1$ year; the risk-free rate is $r = 0.04$; the expected value and volatility of the returns of the risky stock are $\alpha = 0.1$ and $\sigma = 0.3$, respectively; the parameters in the prospect theory utility function (i.e., equation (10)) are calibrated to the estimates of Kahneman and Tversky (1979), as follows: $p = q = 0.88$, $c = 2.25$, and the reference point is set at the investor's initial wealth level, i.e., $R = W_0$; the parameters in the aspiration utility function (i.e., equation (11)) are set as follows: $p = 0.5$, $c_1 = 1.2$, $c_2 = 0$, and the reference point is set at $R = 1.2W_0$.

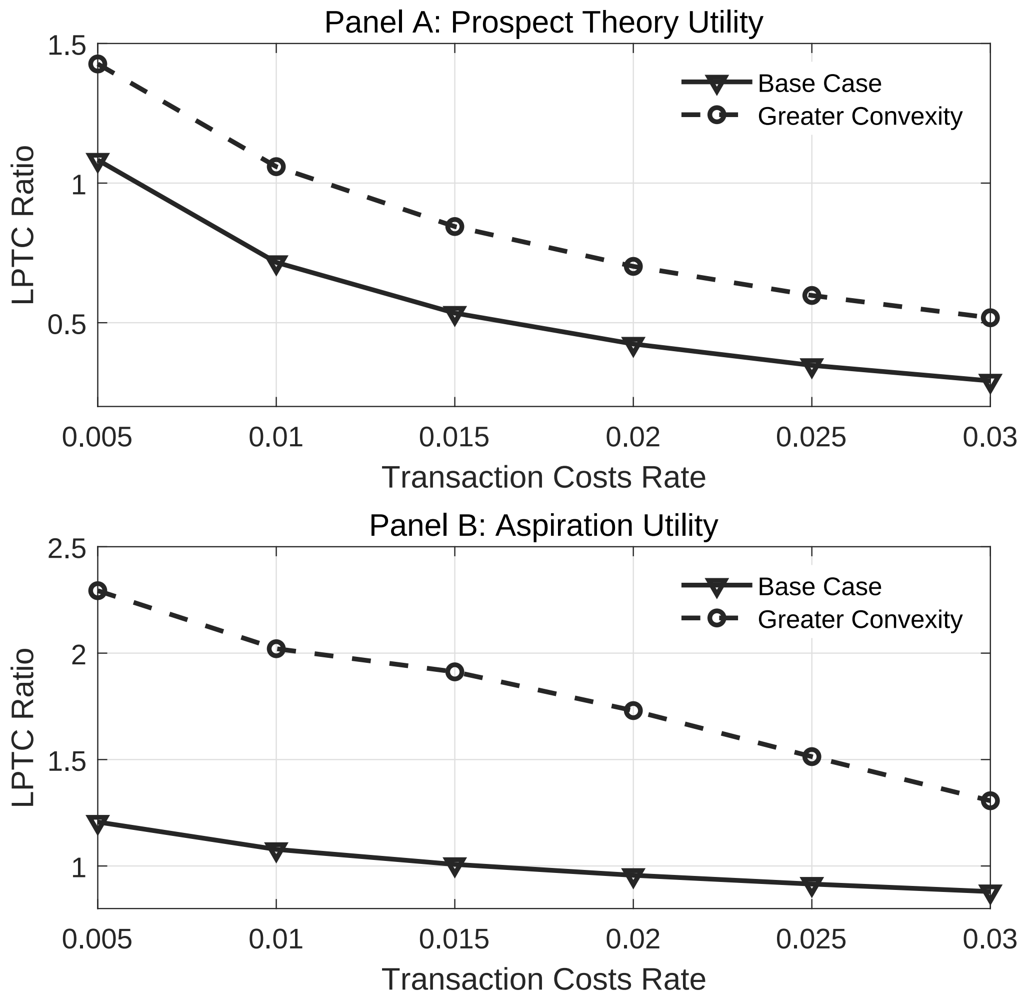

Figure 5: Non-Concave Utility Functions and Liquidity Premia

This figure shows the liquidity premium commanded by the investor when using the prospect theory utility (Panel A) or the aspiration utility (Panel B). The parameter values used to generate these results are as follows: the investor's investment horizon is $T = 1$ year; the risk-free rate is $r = 0.04$; the expected value and volatility of the returns of the risky stock are $\alpha = 0.1$ and $\sigma = 0.3$, respectively; the parameters in the prospect theory utility function (i.e., equation (10)) are calibrated to the estimates of Kahneman and Tversky (1979), as follows: $p = q = 0.88$ ($q = 0.6$ for the case with greater convexity), $c = 2.25$, and the reference point is set at the investor's initial wealth level, i.e., $R = W_0$; the parameters in the aspiration utility function (i.e., equation (11)) are set as follows: $p = 0.5$, $c_1 = 1.2$, $c_2 = 0$, and the reference point is set at $R = 1.2W_0$ ($R = 1.1W_0$ for the case with greater convexity).

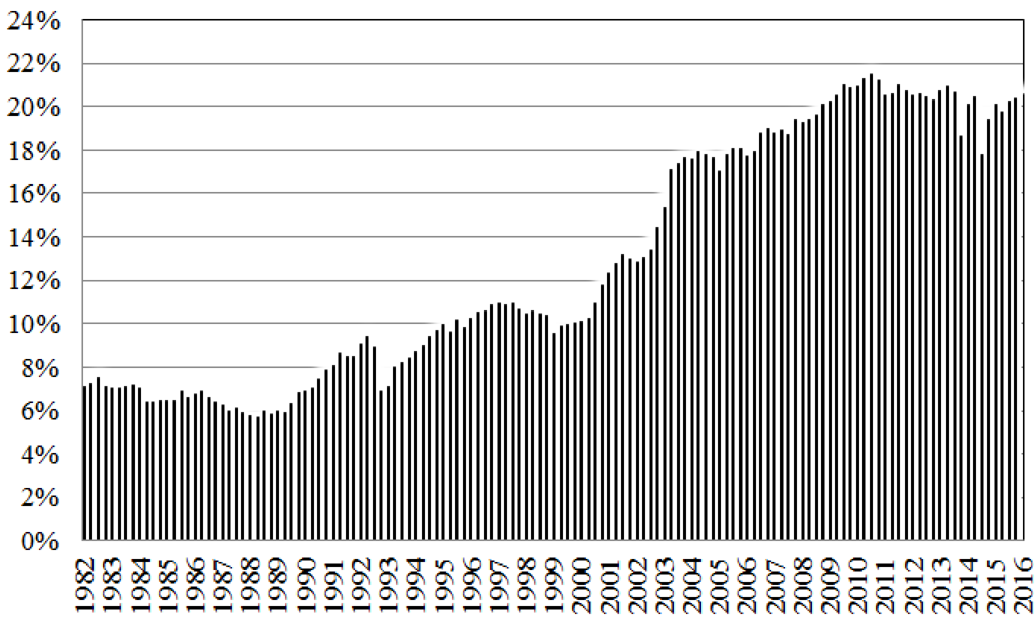

Figure 6: Stock Ownership by Active Mutual Funds

This figure plots the ownership by active mutual funds for each quarter, for the period 1982-2016. For each stock, we take the number of shares held by active mutual funds, and divide it by the total number of shares outstanding. Next, we take the cross-sectional average of this ownership ratio across all stocks in each quarter, and plot the time-series of this cross-sectional average. This only includes stocks listed on the NASDAQ.