NUS RMI Working Paper Series – No. 2021-03

Convex Incentives and Liquidity Premia

Min DAI, Luis GONCALVES-PINTO, Jing XU and Cheng YAN

10 April 2021

NUS Risk Management Institute
21 HENG MUI KENG TERRACE, #04-03 I3 BUILDING, SINGAPORE 119613

www.rmi.nus.edu.sg/research/rmi-working-paper-series

Convex Incentives and Liquidity Premia [*]

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

April 10, 2021

Abstract

We show that convexity in investors' preferences can significantly amplify the effect of transaction costs on the liquidity premia of stocks. This result is derived from the dynamic portfolio problem of fund managers who engage in risk-shifting to capture year-end bonuses, but is robust to other sources of convexity such as loss aversion or status concerns. The larger premia compensate primarily for the lower bonuses resulting from the suboptimal implementation of risk-shifting strategies. Using data on actively-managed mutual funds, we provide empirical support for the novel predictions of our 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, 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, the Chinese University of Hong Kong, and the National University of Singapore.

† 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 convexity in investors' preferences can significantly amplify the effect of transaction costs on the liquidity premia of stocks. This result is derived from the dynamic portfolio problem of fund managers who engage in risk-shifting to capture year-end bonuses, but is robust to other sources of convexity such as loss aversion or status concerns. The larger premia compensate primarily for the lower bonuses resulting from the suboptimal implementation of risk-shifting strategies. Using data on actively-managed mutual funds, we provide empirical support for the novel predictions of our model.

"For actively managed funds, the people that make the ultimate investment decisions are not the owners. If the people making the investment decisions obtain a high reward when things go well and a limited penalty if they go badly they will be willing to pay more than the discounted cash flow for an asset. This is the type of incentive scheme that many financial institutions give to investment managers."

Allen (2001)

1 Introduction

It is empirically recognized that liquidity is valuable to investors and that they demand a return premium to compensate for transaction costs (e.g., Amihud and Mendelson (1986), Eleswarapu (1997), Hasbrouck (2009)). However, few existing theories have been able to corroborate such empirical findings. Constantinides (1986) first argued that, under constant investment opportunities, trading costs significantly reduce the frequency and volume of trading, but have surprisingly little impact on utility. Thus, the return that a marginal investor is willing to exchange for zero trading costs (i.e., the liquidity premium) was found to be an order of magnitude smaller than the transaction cost rate. To increase the magnitude of liquidity premia, Jang, Koo, Liu, and Loewenstein (2007), Lynch and Tan (2011), and Dai, Li, Liu, and Wang (2016) introduce time-variation in investment opportunities.

In this paper, we investigate the effect of investors' convex incentives on the magnitude of liquidity premia. 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 find that the interaction of convex incentives with transaction costs significantly (and endogenously) amplifies the magnitude of liquidity premia to levels comparable with empirical evidence. We establish this result both theoretically and empirically, by focusing on the risk-shifting incentives derived from the year-end bonus that are prevalent in the compensation contracts of mutual fund managers (see Ma, Tang, and Gomez (2019) and Lee, Trzcinka, and Venkatesan (2019)). But we show that this result is robust to other non-concave incentives from the behavioural economics literature, such as those derived from loss aversion (Kahneman and Tversky (1979)) or 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 have become the marginal investors in many stocks.[2] Second, the compensation contracts of portfolio managers typically include a convex component in the form of a year-end bonus (Ma, Tang, and Gomez (2019)).[3] They reward good performance but do not penalize bad performance, giving incentives to take excessive risks as implied by Allen (2001) in the introductory quote. Moreover, these contracts are typically incomplete, giving fund managers the residual control over the assets in their portfolios (Hart (2017)).[4] Third, mutual fund managers that are prone to engaging in excessive risk-taking also tend to hold more illiquid stocks in their portfolios compared to other funds (Huang, Sialm, and Zhang (2011)), suggesting that gambling activities induced by convex incentives are specially likely to interact with stock trading costs. Fourth, data on mutual fund characteristics and portfolio holdings 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

[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] These types of convex incentives are also prevalent in the hedge fund industry. They arise from the typical fee structure that includes a flat management fee plus a performance fee awarded when performance is better than a hurdle rate or a high-water mark. These fee structures also induce excessive risk taking (Lan, Wang, and Yang (2013)), and our argument should hold in this setting as well. We do not focus our analysis on hedge funds because data on mutual funds it more readily available to test our model predictions, which we do in Section 5.

[4] In fact, 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 gambling, which deems these contracts incomplete.

[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.

perfectly liquid, and she can also invest in an illiquid non-benchmark stock which is subject to transaction costs and provides some unspanned risk relative to the benchmark stock. 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. If the fund's portfolio overtakes the benchmark in the interim period, it is optimal for the fund manager to lock-in this relative advantage by replicating the benchmark closely.

An immediate implication of this policy is that it requires high portfolio turnover to resolve the potentially strong portfolio dislocations.[6] Such high turnover requirements imply heavy trading cost bills to be incurred by fund managers. We simulate the optimal policy under our baseline setup and find that the expected discounted value of the trading cost payments in the presence of convex incentives is more than 60 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 prohibitively 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 4.61%.[7]

[6] 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%. However, the implication that benchmark-linked incentives lead to higher turnover are not novel. This result has been obtained by Cuoco and Kaniel (2011) and Sotes-Paladino and Zapatero (2019), for instance. However, their models do not account for the effects of transaction costs.

[7] The manager's risk-adjusted bonus is defined as the minimum amount of additional AUM that the

Therefore, as marginal investors, fund managers demand high liquidity premia to compensate for (i) the heavy trading costs associated with high portfolio turnover, and (ii) the reduction in total compensation for the portfolio manager due to the strong portfolio dislocations that are difficult to undo because of transaction costs. We decompose the liquidity premia into these two parts and show that the latter one is the primary contributor to the amplification effect in our model.

Our results can help reconcile the longstanding disconnect between theory and evidence regarding the magnitude of liquidity premia, from a novel perspective. Amihud and Mendelson (1986) find that the liquidity premium to transaction cost (LPTC) ratio is 1.90 for NYSE stocks, but this contrasts with the model-implied LPTC ratio of 0.07 in Constantinides (1986).[8] In our model, the estimates of liquidity premia are the same order of magnitude as the empirical findings. We use a collar function to model year-end bonuses, and match the moments of the bonus distribution to the empirical estimates provided in Ma, Tang, and Gomez (2019), to find that our model-implied LPTC ratio increases to 1.175.[9] If we remove the convexity from our model, the LPTC ratio drops to 0.02, similar to the result in Constantinides (1986).[10]

In prior work, Jang, Koo, Liu, and Loewenstein (2007) and Lynch and Tan (2011) introduce time-variation in investment opportunities to increase (exogenously) the frequency and volume of trading for investors, allowing them to generate (mechanically) larger liquidity premia, but not large enough to match the empirical findings. More recently, Dai, Li, Liu,

---

manager requires for waiving her bonuses. This measure is a function of managerial risk aversion. The result is computed under the assumption that trading costs are waived upon trading, to prevent the mechanical effect that trading costs can have on relative performance. It is important to highlight that, even though we waive the trading cost payments, the optimal policy that is followed is a policy that includes the no-trading region generated by those transaction costs.

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

[9] 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 used in Basak, Pavlova, and Shapiro (2007). These additional results are available from the authors upon request.

[10] 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).

and Wang (2016) show that one needs very frequent changes in the investment opportunity set to generate theoretical liquidity premia that are comparable in magnitude to the empirical estimates. In contrast, the large premia in our model are generated endogenously, as we keep investment opportunities constant like in Constantinides (1986). The strong portfolio distortions and high turnover in our model are the result of the optimal response to the convexity of the bonus function.

The main empirical prediction of our model is that, the stronger the risk-shifting incentives induced by the convexity in the compensation structure, the larger the LPTC ratio. We use a sample of U.S. domestic actively-managed equity mutual funds and their portfolio holdings to provide support for this prediction. Specifically, we test whether stronger risk-shifting incentives in the fund industry are associated with higher liquidity premia. 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. We use the eight risk-shifting measures proposed by Huang, Sialm, and Zhang (2011). 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 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 component. The final result is a single stock-level proxy that captures the convex incentives of the mutual funds that hold the stock.[11] We then examine how this proxy for stock-level

[11] 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.

convex incentives affects the relation between transaction costs and future stock returns.

We regress excess returns on lagged effective trading cost estimates and their interaction with convex incentives. In our list of control variables we include stock characteristics such as size, turnover, and ownership by active mutual funds, among others. We start by confirming the findings in Hasbrouck (2009) that effective trading costs are strongly related to future stock returns. But 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 incentives require more than 1% higher excess return per month. These results are consistent with our theory, that convex 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 theory 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 lost compensation. In the data, we find that the average turnover of the portfolios of stocks held by funds with strong convex incentives is about double that in the weak incentive group. The difference is especially large when stocks have low betas, going from double to three times larger. However, for stocks with high betas, turnover is only about 40% larger in the strong incentive group. This is consistent with the conjecture in Boguth and Simutin (2018) that low beta stocks provide lower leverage, and gambling 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 gamble using liquid stocks, requiring additional future return to gamble with illiquid stocks.

The remainder of this paper unfolds as follows. Section 2 reviews the related literature. Section 3 presents the theoretical framework. Section 4 describes the numerical analysis of the optimal investment policy, the magnitude of the liquidity premia, and their sensitivity to changes in parameter values. Section 5 provides empirical evidence to support the novel predictions of the theoretical model. Section 6 concludes. We relegate to the Appendix all the technical issues, additional results, and details on the construction of the empirical variables.

2 Related Literature

This paper is related to the intersection of two strands of research. First, the research on the impact of transaction costs on portfolio choice and liquidity premia. Second, the research on the risk-shifting incentives derived from contracts with option-like characteristics.

The seminal work of Constantinides (1986) shows that introducing transaction costs into the portfolio choice problem of Merton (1969) leads to a drastic reduction in the frequency and volume of trading. More importantly, it shows that the expected return that the investor is willing to exchange for zero trading costs is surprisingly small relative to the transaction cost rate. The author concludes that transaction costs only have a second-order effect on liquidity premia.

However, this conclusion is puzzling because it is not in line with many empirical findings that suggest that transaction costs significantly affect the time-series and the cross-section of expected stock returns, such as Amihud and Mendelson (1986).

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, there exist three main references. First, is the paper by Jang, Koo, Liu, and Loewenstein (2007). They argue that the main reason for the puzzling disconnect is the assumption of constant investment opportunities in Constantinides (1986). They show that, by extending the model to include time-varying investment opportunities, in the form of stock market regime shifts, transaction costs can have a larger effect on liquidity premia. The main driver of their results is the increased amount and frequency of trading induced by the exogenous (and fully observable) market shifts. Yet, the liquidity premium to transaction cost (LPTC) ratio that they find, for a reasonable calibration of their regime-switching model, is only 0.25, for a proportional transaction cost rate of 1%. This figure is not large enough to fully explain the liquidity premium puzzle.[12]

The second main reference is Lynch and Tan (2011). They use a discrete-time framework

[12] In a recent paper by Chen, Dai, Goncalves-Pinto, Xu, and Yan (2020), they relax the assumption used in Jang, Koo, Liu, and Loewenstein (2007) that regime shifts are fully observable. When regime shifts are unobservable, and investors need to infer the current market regime from past price movements, the model-implied LPTC ratio increases significantly. They show that the main driver of their results is the suboptimal risk exposure, as turnover is low when investors cannot observe the regime shifts.

and incorporate return predictability, labor income, and state-dependent transaction costs. Their model generates an LPTC ratio of 0.43, for a proportional transaction cost rate of 2%. This is also significantly lower than the empirical estimates.

More recently, Dai, Li, Liu, and Wang (2016) consider the possibility that markets can close for trading. However, this is not the feature that drives their results. Instead, it is the fact that they assume that stock return volatility changes between trading and non-trading periods. This is equivalent to the assumption of time-varying investment opportunities. This model is quite successful at bridging the gap between theory and evidence regarding the magnitude of liquidity premia. For instance, when they take the volatility in trading periods to be three times higher than that in non-trading periods, the LPTC ratio that they obtain is 1.76, for a proportional transaction cost of 1%.

It is important to highlight that the common ingredient in the three references described above is the (exogenous) time-varying nature of investment opportunities. The intuition for their results is quite simple. If the investment opportunities are changing over time, then an investor who faces no transaction costs will adjust her asset allocations immediately to the conditions of the new regime. When transaction costs are introduced, the investor either trades very frequently and pays a heavy transaction cost bill (Jang, Koo, Liu, and Loewenstein (2007) and Lynch and Tan (2011)), or she does not trade as frequently but loses utility from the highly suboptimal risk exposure (Dai, Li, Liu, and Wang (2016)).

We contribute to this theoretical research by proposing an alternative mechanism that amplifies the effect of transaction costs on liquidity premia. We keep investment opportunities constant, but offer an incomplete contract with option-like characteristics to the individual who has residual control over the assets in the portfolio (i.e., the portfolio manager). These conditions are common in the mutual fund industry. In this industry, households typically delegate their investment decisions to an investment advisor, who in turn hires a fund manager. The contract between the investment advisor and the fund manager typically includes an option-like component associated with the performance of the fund relative to a self-designated market index, and this component is generally asymmetric: managers receive a bonus at year-end if they outperform the index, but are not penalized in case of underperformance. Farnsworth and Taylor (2006) and Ma, Tang, and Gomez (2019) document that the inclusion of year-end bonuses in compensation packages is pervasive in the mutual fund industry. Lee, Trzcinka, and Venkatesan (2019) show that risk-shifting of mutual fund managers is motivated more by this type of compensation structure than by a tournament to capture flows.[13]

3 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_{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

[13] 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, a few studies have shown that the tournament effect disappears when using different risk-shifting 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-shifting 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.

then have the following sub-wealth processes for liquid and illiquid portfolio holdings:

$$dX_t = (rX_t + \xi_t(\alpha_1 - r))dt + \sigma_1 \xi_t dW_{1t} - (1+\lambda)dL_t + (1-\mu)dM_t \quad (2)$$

$$dY_t = \alpha_2 Y_t dt + \sigma_2 Y_t dW_{2t} + dL_t - dM_t \quad (3)$$

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 time T consists of two parts: a management fee proportional to the value of assets under management (AUM), and a performance-based bonus that is determined by the fund's performance relative to an external benchmark portfolio. 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:

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

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:

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

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 for the fund and the benchmark, respectively, over the period $(0,T)$. Since only the growth rate of the fund's portfolio over that of the benchmark matters for the calculation of the bonus, we can set $Z_0 = X_0 + Y_0$ without loss of generality.

Consistent with empirical evidence, we consider a bonus function of the collar-type:

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

where $f_L \geq 0$, $\psi = (f_H - f_L)/(\theta_H - \theta_L) > 0$, 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]

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

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

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]

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. Thus, we include 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.

4 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

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

[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. $\beta = 1$. We assume a transaction cost rate of 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 $f_H = 1.5$. This means that, when the performance of the fund's portfolio in excess of that of the benchmark is below 1% annually, the fund manager receives no bonus. If instead the fund's portfolio is outperforming 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.

4.1 Optimal Stock Allocations

In our model, the optimal policy is a function of the performance of the fund relative to the benchmark. 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 overweighting the non-benchmark stock in the portfolio and reducing the exposure to the benchmark stock. This increases 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 non-benchmark stock and tracking 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 like in Merton (1969) for the case without trading costs 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, which is similar to the optimal policy 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 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.

[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 4.2, which is the central focus of our paper.

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 resembles the policy in Panel B for extreme values of relative performance, which is flat and 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, when convex incentives are absent, and the transaction cost rate is 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 bonuses, 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. We study this question in more detail in the next section.

4.2 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 dislocations become hard to undo because of transaction costs. We also examine the sensitivity of liquidity premia to changes in the values of the input parameters.

4.2.1 Definition and Decomposition

Assume that we have two stocks that are perfectly correlated and have the same volatility, but one of them is perfectly liquid and the other is subject to transaction costs. For both of these stocks to be held in equilibrium, the expected return of the stock that is subject to transaction costs must exceed that of the liquid one. 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 fashion: it is the quantity δ that solves the following equation:[18]

$$\varphi(0, \zeta^*, 0; \mu, \lambda, \alpha_2) = \varphi(0, \zeta^*, 0; 0, 0, \alpha_2 - \delta) \tag{8}$$

where the function $\varphi(t, \zeta, \eta; \mu, \lambda, \alpha_2)$ is as defined in (23) in Appendix A, and

$$\zeta^* = \arg\max_{\zeta \in [0,1]} \varphi(0, \zeta, 0; \mu, \lambda, \alpha_2) \tag{9}$$

is the optimal initial allocation to the illiquid non-benchmark stock.[19]

Intuitively, transaction costs decrease the manager's utility through two channels. On the one hand, the payment of transaction costs directly reduces the amount of AUM. On the other hand, the fund manager is forced to have a risk exposure that is suboptimal, compared to the case without transaction costs. In order to differentiate these two effects, we solve the following equation (like in Dai, Li, Liu, and Wang (2016)):

$$\varphi^A(0, \zeta^*, 0; 0, 0, \alpha_2) = \varphi(0, \zeta^*, 0; 0, 0, \alpha_2 - \delta^0)$$

[18] This is equivalent to the following equation expressed in terms of the original value function: $V(0, x, y, z; \mu, \lambda, \alpha_2) = V(0, x, y, z; 0, 0, \alpha_2 - \delta)$, subject to $z = x + y$, $y/(x + y) = \zeta^*$, where the value function V is defined in Appendix A.

[19] In the definition of liquidity premium, we set $\eta = 0$ because at the start of the investment period the manager's portfolio and the benchmark portfolio have the same value, i.e., $Z_0 = X_0 + Y_0$. In addition, we choose the optimal ζ in the value function, hence implicitly assume that the fund manager can choose her optimal initial position at zero cost. Relaxing this assumption does not change our main results. It can only lead to even higher liquidity premia, because the manager would have to make a lump-sum purchase at the initial time at a cost.

where φ^A is the value function when we assume the manager's optimal strategy (ξ_t^*, L_t^*, M_t^*) is implemented but transaction costs are waived. Therefore, δ^0/δ measures the fraction of the liquidity premium that is due to portfolio displacement compared to trading cost payments.

Table 2 reports the characteristics of the optimal trading policy (at time $t = 0$) and the liquidity premia for various levels of transaction costs. In order to facilitate comparison, we first present in Panel A the case in which the bonus is absent. 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 only 0.017 (0.009). If the fund manager is the marginal investor in this stock, her trading on this stock commands 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 compared to the case without bonuses (δ_c).

Panel B also reports the maximum and minimum values for the buy and sell boundaries. For instance, if the trading cost rate is 1%, the buy boundary goes from a minimum $B_*(0, \eta) = 0.00$ 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 for the fund manager 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 of trading costs and to portfolio displacements. For example, for a 1% transaction cost rate, δ^0/δ equals 94.68% in the absence of bonuses (Panel A), and equals 50.77% in the presence of bonuses (Panel B). This means that direct payments contribute more to liquidity premia in the presence of bonuses, which is consistent with the results reported in Table 1.[20]

The results in Panel B suggest that the costs due to portfolio displacement are also 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 (ξ_t^*, L_t^*, M_t^*) under transaction costs but these costs can be 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%, even when the trading costs are assumed to be waived.

[Insert Figure 2 about here]

4.2.2 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. In the following, we briefly discuss some of these results.

[Insert Table 3 about here]

Return correlation. The correlation between the returns of the benchmark and non-benchmark stocks is an important determinant of the effectiveness of risk-shifting. Intuitively, assets which exhibit low correlation with the benchmark are better tools for gambling purposes. If the correlation increases, then the benefit from risk-shifting decreases and

[20] The ratio δ_c^0/δ_c in our model is larger than in Dai, Li, Liu, and Wang (2016) for two reasons. First, we consider a relatively short investment horizon of 1 year, and the investor does not trade when near maturity. Thus, the cost due to portfolio displacement is larger. In contrast, they assume a long horizon, and this effect is weaker in their model. Second, our investor can trade a perfectly liquid stock which is positively correlated with the illiquid stock, which further reduces her demand for trading the illiquid stock.

the fund manager gambles less with the illiquid stock. This is then expected to reduce the liquidity premia. Table 3 suggests that it 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. In order to outperform a riskier benchmark, the manager needs to overweight the illiquid non-benchmark stock even more, as it provides some risk exposure unspanned by the benchmark stock. This in turn leads to larger liquidity premia. Table 3 shows the changes in 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, and it increases to 1.175 when the benchmark portfolio is fully invested in liquid stock (i.e., $\beta = 1$, 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 should 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 and the risk-shifting incentives that result from it. This is the main prediction tested in our empirical analysis of Section 5.

4.2.3 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 gambling 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. Trading on the illiquid non-benchmark stock is more likely to happen if the position limits on the liquid stock are binding. Therefore, one could expect that, if such position limits are relaxed, the demand for illiquid stock would decrease, and the impact of trading costs on the derived utility of the fund manager would weaken.

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 $[0, 1]$, but we relax the position limit on the liquid benchmark stock to lie in the interval $[-a, 1 + a]$, where a ranges from 0 to 2.[21] The results show that the LPTC ratio only decreases slightly for looser position limits on the liquid benchmark stock.

[Insert Figure 3 about here]

4.3 Other Sources of Convexity: Reference-Dependent Utility

In our main model, the utility function of the economic agent (the fund manager) is globally concave, and the agent's convex incentives stem from the convexity of the 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. For example, 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, and that the effect is not restricted to the particular model that we examine in the baseline setup of this paper.

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

$$E[U(W_T; R)] \tag{10}$$

where W_T is the investor's gross wealth level at time T, and R is a reference point which

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

determines the investor's utility level.[22]

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

$$U(W;R) = \begin{cases} (W-R)^p & \text{if } W \geq R \\ -c(R-W)^q & \text{if } W < R \end{cases} \tag{11}$$

where $0 < p, q < 1$ and $c > 0$.

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

$$U(W;R) = \begin{cases} \frac{W^p}{p} & \text{if } W < R \\ c_1 \frac{W^p}{p} + c_2 & \text{if } W \geq R \end{cases} \tag{12}$$

where $0 < p < 1$, $c_1 \geq 1$ and $c_2 \geq 0$. This aspiration utility specification captures the idea that, besides normal consumption, the investor also cares about the 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 are set 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 (11) are calibrated to the estimate 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 parameters in function (12) are set as follows: $p = 0.5$, $c_1 = 1.2$, $c_2 = 0$, and the reference point is set at $R = 1.2W_0$.[23]

Figure 4 plots the investor's optimal allocation to the risky stock, as a function of her

[22] Like what we do in our main model, we assume the investor derives utility from gross wealth level at time T to exclude the mechanical effect of transaction costs on wealth upon liquidation.

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

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 allocation in response to fluctuations in wealth level. 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 rapidly, implying a strong gambling incentive. Therefore, similar to the intuition developed in our main model, it can be expected that the presence of stock transaction costs will be particularly burdensome to the investor in this setting.

[Insert Figure 4 about here]

Figure 5 plots the LPTC ratio against the transaction costs rate. The liquidity premium is calculated at the initial time. It suggests that the liquidity premium can be substantial when the investor exhibits either prospect theory utility or aspiration utility. For example, with a transaction costs rate of 1%, the LPTC ratio is above 0.7 in the prospect theory utility case, and above 1.0 in the aspiration utility case. Figure 5 also shows the LPTC ratio when we increase the convexity of the 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 could significantly amplify the magnitude of liquidity premia. These results are consistent with the results that we derived from our baseline model in Section 3.

[Insert Figure 5 about here]

5 Empirical Analysis

In this section, we provide evidence in support of the main implications of our theory model. First, we show that the incentive of mutual fund managers to engage in risk-shifting leads to higher turnover for the stocks in their portfolios. Second, we show that trading costs have a positive relation with expected stock returns, especially for stocks held by mutual funds with strong risk-shifting incentives. To the best of our knowledge, the existing empirical evidence has not yet verified any of these novel implications of our model.

5.1 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 stock positions of funds (i.e., the number of shares of a stock held by a given mutual fund) and provides identifiers for fund families (Wermers (2000)).[24]

The latest update of the MFLINKS file on WRDS includes the mapping of mutual funds to their portfolio holdings from 1980 to 2016. This allows us to set 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).[25] In addition, in 2004 the CRSP mutual fund database switched its data provider from Morningstar to Lipper, and the overlap between these two versions of the database is imperfect.

We restrict our analysis to diversified domestic actively-managed equity mutual funds.[26] We compute fund-level variables by aggregating across all the share classes. For instance, the TNA of the fund is the sum of the TNAs of all its share classes.

[24] TFMFH only reports relatively large portfolio positions, 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.

[25] 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 looking at the ratio between liquidity premium and trading cost rate, and this ratio can still be large, even if the liquidity premium is small.

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

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 of less than $5 million, because Elton, Gruber, and Blake (2001) show that the returns of such small funds tend to be biased upwardly in the CRSP database.[27]

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 from 2004 to 2016 for most of our main empirical tests, because it is more likely that these funds will be marginal investors in these stocks during this period.

[Insert Figure 6 about here]

5.2 Main Variables

The main prediction of our theory model is that risk-shifting in the mutual fund industry is an important factor driving the relation between transaction costs and expected stock returns. 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). The data is available in Joel Hasbrouck's website, but it ends in 2009.[28] We extend his dataset

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

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

until 2017, following the same estimation method. Hasbrouck (2009) performs separate analyses for NASDAQ, AMEX, and NYSE stocks. We focus part of our analysis on NASDAQ stocks, like in Eleswarapu (1997), but our results remain qualitatively similar using stocks from the other two exchanges. After merging all the datasets of the previous section with this dataset on effective trading costs, our final sample includes 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 propose eight measures of risk-shifting. We provide a detailed description of each of these measures in Appendix C. We first calculate each of these fund-level proxies for every mutual fund in our sample. Next, we aggregate each proxy across all funds holding a given stock, using their quarterly share holdings as weights. This gives us one stock-level measure for each proxy. We formulate these proxies such that large values correspond to greater propensity to risk-shifting, and therefore stronger incentives by the holders of the stock to engage in gambling.

We then create a composite measure at the stock-level. We use two methods to do this. First, we use the average percentile rank across the eight proxies. In December of the prior year, we assign percentile ranks to each stock-level proxy for the whole cross-section of stocks. Then, we compute the average percentile rank across the proxies 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 stock-level proxies, and use the first principal component obtained from that analysis. We denote this composite measure as FPC.[29]

[29] 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 proxies instead of the composite measures in our empirical analysis, and the results are qualitatively similar. These results are available from the authors upon request.

5.3 Portfolio Statistics and Empirical Specification

In this section, we test how risk-shifting incentives (RSI henceforth) of stock holders affect the relation between effective trading costs and expected 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). 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 the 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 of Hasbrouck (2009). In total, this three-way sort results in 50 portfolios, and the number of stocks per portfolio is around 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 theory model, we show that convex incentives increase stock turnover. This is indeed what we find in the data. Table 4 shows that, when we compute RSI using APR in Panel B (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]

We highlight that, the difference in turnover is specially 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 an average turnover that is about 150% larger in the above-median RSI group, compared with the below-median 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 RSI group. The difference is similar in Panel B. This is consistent with the idea in Boguth and Simutin (2018) that low beta stocks provide lower leverage, and gambling with such stocks requires additional trading.

Next, we examine the relation between trading costs and expected stock returns, and especially how this relation is affected by convex incentives in the mutual fund industry. We use ex-post returns as a proxy for expected returns. We regress the expected returns (in excess of the risk-free rate) of the 50 portfolios on the lagged effective trading cost measure of Hasbrouck (2009) and its interaction with an indicator variable for above-median RSI. 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:

$$\begin{aligned} R_{i,t+1} =& \gamma_0 + \gamma_c c_{i,t} + \gamma_d DummyRSI_{i,t} + \gamma_{cd} \ \left(c_{i,t} \times DummyRSI_{i,t}\right) \\ &+ \gamma_t Turnover_{i,t} + \gamma_{ct} \ \left(c_{i,t} \times Turnover_{i,t}\right) \\ &+ \gamma_l LRMC_{i,t} + + \gamma_o Ownership_{i,t} \\ &+ \gamma_m \beta_i^m + \gamma_s \beta_i^{smb} + \gamma_h \beta_i^{hml} + \epsilon \end{aligned} \tag{13}$$

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). $Turnover$ is the average ratio of trading volume to shares outstanding across the stocks in each portfolio. $Ownership$ is the ratio of the number of shares held by active mutual funds to the total number of shares outstanding, for a given stock. The effective trading cost measure is denoted as c. To assess the impact of RSI on the relation between trading costs and future excess returns, the regression specification includes an indicator function $DummyRSI$, which is equals one for above-median RSI portfolios, and equals 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.

5.4 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 tradig cost variable (i.e., c).

[Insert Table 5 about here]

The difference between Panels A and B is the variable used in the sequential sorting to create the 50 portfolios used in the analysis. In Panel A we use the average percentile rank (APR) across the eight proxies for risk-shifting, and in Panel B we use the first principal component (FPC). This explains the slight differences in the regression coefficients for the trading costs variable across the two panels. These tests control for the variables used in Hasbrouck (2009), which include the relative stock size variable $LRMC$, and the unconditional betas from a Fama-French model.

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

In columns (3) and (8) we include the interaction between trading costs and our indicator $DummyRSI$ to find that it is statistically significant and economically large. For instance, in column (3) of Panel A the interaction between trading costs and the indicator variable suggest 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 qualitatively similar in Panel B where FPC is used.

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 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 our indicator variable with stock turnover, which is the stock trading volume as a percentage of the number of shares outstanding.[30] 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 expected 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 in which 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 interaction between trading cost and the indicator variable has a value of 1.07, which is statistically significant at 1% level. This is equivalent to say 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 group (i.e., the coefficient on c, when $DummyRSI = 0$), which is an economically large effect.

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 theory model.

5.5 Convex Incentives and Turnover

In the decomposition analysis of Section 4.2.1, we show that a significant portion of liquidity premia 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 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 negative relation between trading costs and turnover to be stronger

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

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 than for less liquid ones. In other words, to use illiquid stocks in their gambling activities, fund managers require additional compensation in the form of higher future stock returns.

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

$$\begin{aligned} Turn_{i,t+1} = & \gamma_0 + \gamma_c\, c_{i,t} + \gamma_d \;\; DummyRSI_{i,t} + \gamma_{cd} \;\; \left(c_{i,t} \times DummyRSI_{i,t}\right) \\ & + \gamma_s Ln(Size)_{i,t} + \gamma_a Alpha_{i,t} + \gamma_b Beta_{i,t} + \gamma_v IdioVol_{i,t} \\ & + \gamma_d DivYield_{i,t} + \gamma_o Ownership_{i,t} + \epsilon_{i,t} \end{aligned} \tag{14}$$

where $Turn$ 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 not inconsistent with our main argument. 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.[31]

6 Conclusion

In the mutual fund industry, fund managers' contracts are typically incomplete and include option-like components, which induce fund managers to engage in gambling, as they hold a residual control over the assets in their portfolios. We use this industry as a laboratory to study how convex incentives created by management contracts affect the liquidity premia of stocks.

Theoretically, we show that the optimal response to the convexities embedded in 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 make it more difficult for fund managers to adjust their portfolio positions, which makes gambling less effective at capturing year-end bonuses. We show that the suboptimal risk-shifting is the main contributor to the amplification of the effect of trading costs on liquidity premia. In other words, the aggregate of fund managers, acting as the marginal investor in a given stock, demand high liquidity premia to compensate for the bonuses that are lost due to limitations imposed by trading costs.

Empirically, we show that the interaction between trading costs and convex incentives are significant and economically large. We show that this is not driven by increased turnover,

[31] 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).

implying that they must be driven by portfolio distortions that are hard to undo because of trading costs, hampering fund managers' ability to capture year-end bonuses. Therefore, they require higher future returns to be willing to gamble 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.

Allen, F., 2001, "Presidential Address: Do Financial Institutions Matter?," *Journal of Finance*, 56, 1165–1175.

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.

Amihud, Y., and H. Mendelson, 1986, "Asset Pricing and the Bid-Ask Spread," *Journal of Financial Economics*, 17, 223–249.

Anton, M., and C. Polk, 2014, "Connected Stocks," *Journal of Finance*, 69, 1099–1127.

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.

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

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.

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, 2020, "Incomplete Information and the Liquidity Premium Puzzle," *Management Science*, forthcoming.

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

Farnsworth, H., and J. Taylor, 2006, "Evidence on the Compensation of Portfolio Managers," *Journal of Financial Research*, 29, 305–324.

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.

Kisin, R., 2011, "The Impact of Mutual Fund Ownership on Corporate Investment: Evidence from a Natural Experiment," Working Paper, Washington University in St. Louis.

Lan, Y., N. Wang, and J. Yang, 2013, "The Economics of Hedge Funds," *Journal of Financial Economics*, 110, 300–323.

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.

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 5.

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,

$$V(t,x,y,z) = \max_{\Theta_{[t,T]}} \; E\left[\frac{[(1+f)(X_T+Y_T)]^{1-\gamma}}{1-\gamma} | X_t = x, Y_t = y, Z_t = z\right] \tag{15}$$

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

$$\min\Big\{ -\partial_t V - \mathcal{L}V, \;\; \partial_y V - (1-\mu)\partial_x V, \;\; (1+\lambda)\partial_x V - \partial_y V \Big\} = 0 \tag{16}$$

with terminal condition

$$V(T,x,y,z) = \frac{1}{1-\gamma}\left[\left(1 + f\left(\ln\frac{x+y}{z}\right)\right)(x+y)\right]^{1-\gamma} \tag{17}$$

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:

$$\begin{aligned}\mathcal{L}V &= rx\partial_x V + \alpha_2 y\partial_y V + (r+\beta(\alpha_1 - r))z\partial_z V + \frac{1}{2}\sigma_2^2 y^2 \partial_{yy} V + \frac{1}{2}\sigma_1^2\beta^2 z^2 \partial_{zz} V + \beta\rho\sigma_1\sigma_2 yz\partial_{yz} V \\ &+ \sup_{0\leq\pi\leq 1}\Big\{ \big[(\alpha_1 - r)x\partial_x V + \beta\sigma_1^2 xz\partial_{xz} V + \rho\sigma_1\sigma_2 xy\partial_{xy} V\big]\pi + \frac{1}{2}\sigma_1^2 x^2 \partial_{xx} V\pi^2\Big\} \end{aligned}\tag{18}$$

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:

$$\min \{-\partial_t V - \mathcal{L}V, \;\; \partial_y V - (1-\mu)\partial_x V\} \,|_{x=0} = 0 \tag{19}$$

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:

$$\min \{-\partial_t V - \mathcal{L}V, \;\; (1+\lambda)\partial_x V - \partial_y V\} \,|_{y=0} = 0 \tag{20}$$

The homogeneity of the CRRA preferences, and the linearity of (2), (3), (4) and (5) from section 3, 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:

$$V\left(t, \frac{x}{x+y}, \frac{y}{x+y}, \frac{z}{x+y}\right) = \frac{1}{(x+y)^{1-\gamma}} V(t, x, y, z) \tag{21}$$

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 (21) by the following:

$$h(t, \zeta, \eta) = V(t, 1-\zeta, \zeta, e^{-\eta}) \tag{22}$$

and further define

$$\varphi(t, \zeta, \eta) = \frac{1}{1-\gamma} \log[(1-\gamma) h(t, \zeta, \eta)] \tag{23}$$

which leads to

$$V(t, x, y, z) = \frac{(x+y)^{1-\gamma}}{1-\gamma} e^{(1-\gamma)\varphi(t,\zeta,\eta)} = \frac{[(x+y)e^{\varphi(t,\zeta,\eta)}]^{1-\gamma}}{1-\gamma} \tag{24}$$

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:

$$\min\{-\varphi_t - \mathcal{M}\varphi, \;\; \mathcal{M}_1\varphi, \;\; \mathcal{M}_2\varphi\} = 0 \tag{25}$$

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

$$\varphi(T,\zeta,\eta) = \ln\left(1 + f(\eta)\right). \tag{26}$$

Given the solution to equation (25) with terminal condition (26), 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 (25), we consider the following penalty approximation:

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

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

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

Step 3: Solve the discretized version of the following equation:[32]

$$\varphi_t^{i+1} + \mathcal{M}_0(\pi_i^*)\varphi^{i+1} + K(-\mathcal{M}_1\varphi^{i+1})^+ + K(-\mathcal{M}_2\varphi^{i+1})^+ = 0, \quad (28)$$

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.

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

The fund managers would 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), which studies 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 recently 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 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 results using this method are reported in Panel A. 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. The results using this method are reported in Panel B. 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, while *MKT Beta*, *SMB Beta*, and *HML Beta* are the unconditional betas obtained from a three-factor Fama-French model estimated over the entire sample period for each portfolio. *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					Panel B: First Principal Component				
	(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 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 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

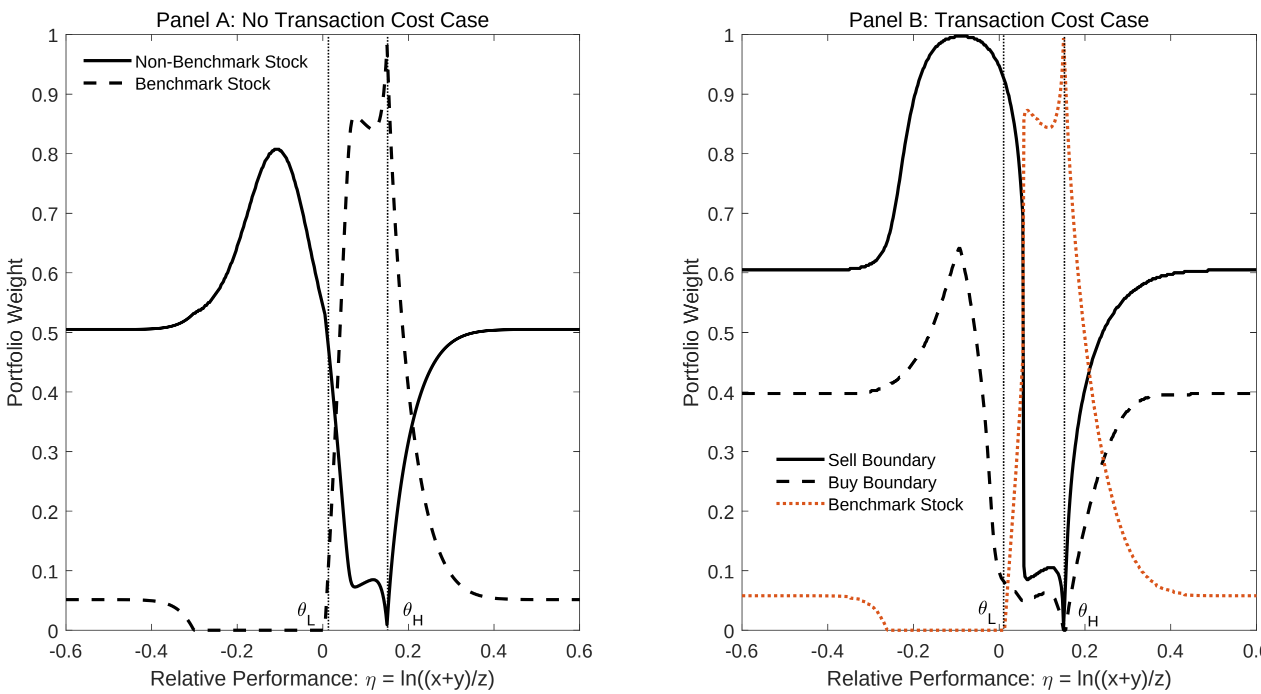

Figure 1: Optimal Stock Allocations

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 $f_H = 1.50$.

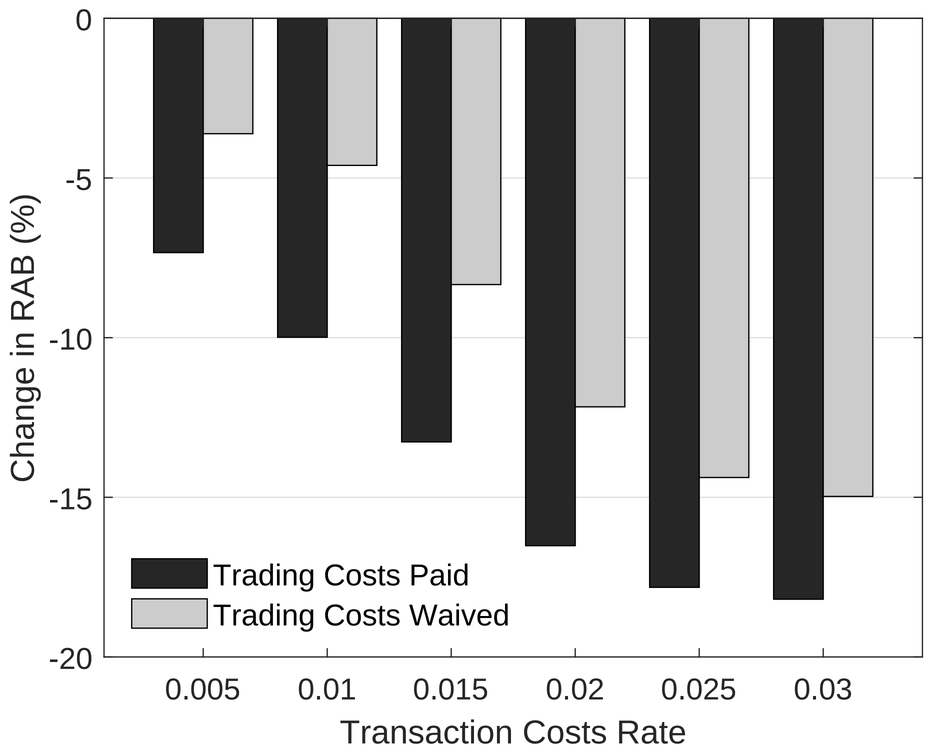

Figure 2: Trading Costs and Reduction in Risk-Adjusted Bonuses

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 $f_H = 1.5$.

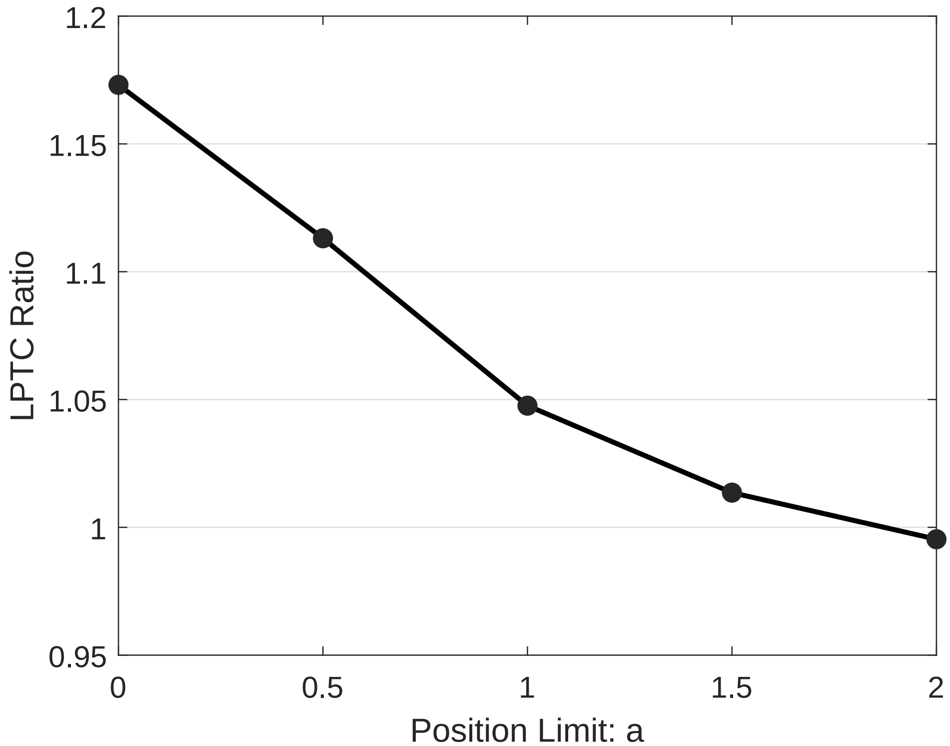

Figure 3: Relaxing the Position Limits on the Liquid Benchmark Stock

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 $f_H = 1.5$.

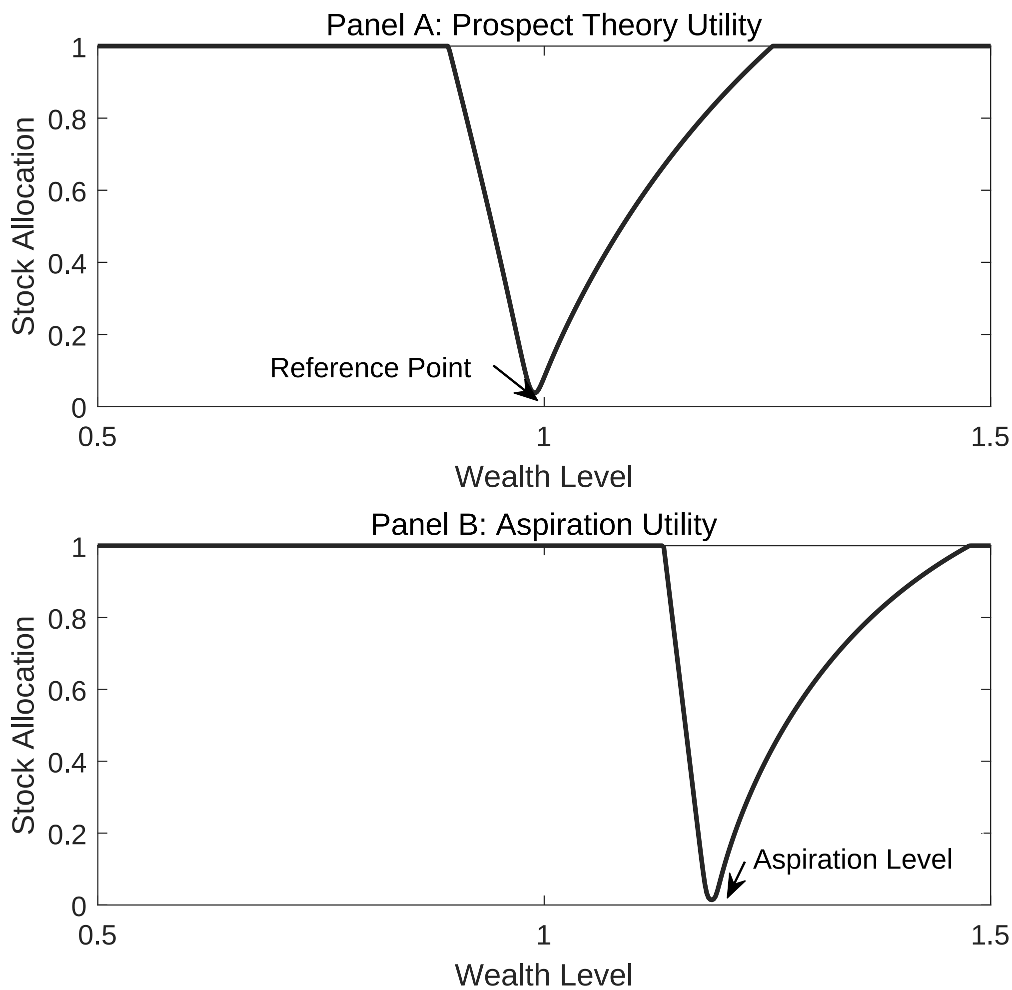

Figure 4: Optimal Trading Strategy with Non-Concave Utility Functions

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 (11)) 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 (12)) are set as follows: $p = 0.5$, $c_1 = 1.2$, $c_2 = 0$, and the reference point is set at $R = 1.2W_0$.

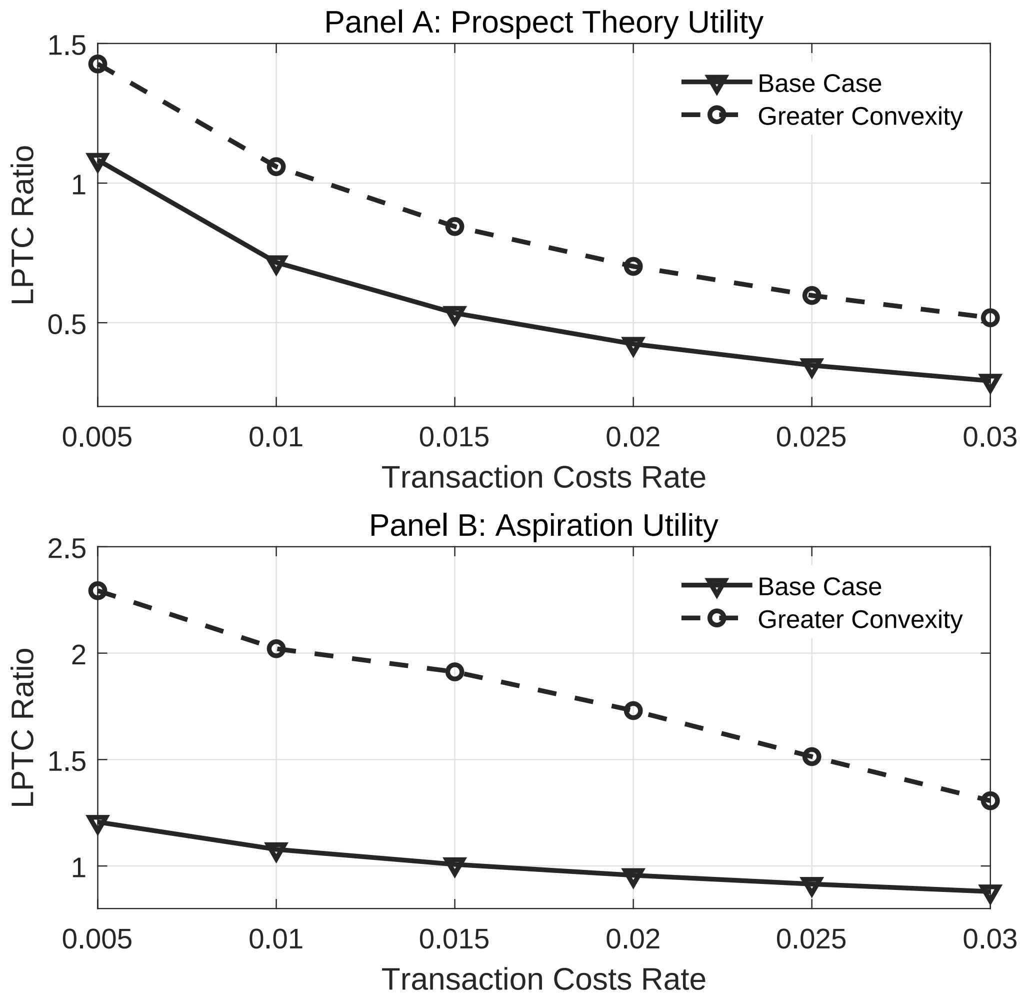

Figure 5: Non-Concave Utility Functions and Liquidity Premia

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 (11)) 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 (12)) 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).

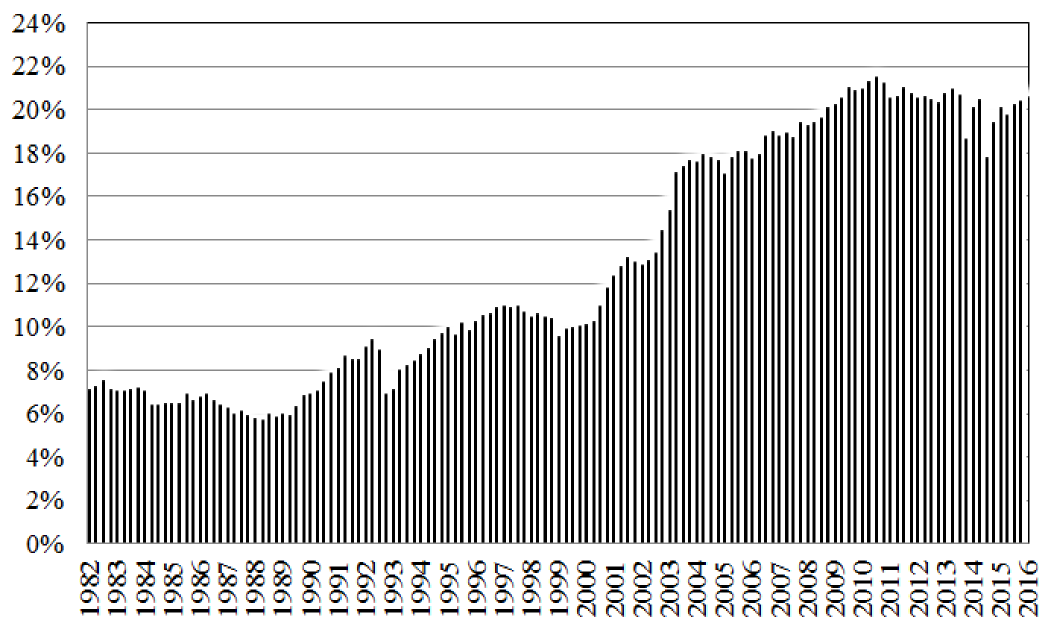

Figure 6: Stock Ownership by Active Mutual Funds

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.