Delegated Portfolio Management and Risk-Taking of Hedge Funds

Jens Jackwerth

August 7, 2020

Abstract

When an investor delegates portfolio management to a hedge fund manager, whose risk-taking preference governs? Single-period models suggest stark variation in risk-taking across fund value and time as fund managers maximized their own well-being. Empirical validation is hard to come by, as each hedge fund traces out only a few points on that risk-taking surface. Cross-sectional pooling of normalized returns allows precise estimation of the normalized risk-taking surface. In fact, it is almost flat with some increased risk-taking at low fund values. A multi-year model is consistent with the findings.

Keywords: Risk-taking, hedge funds

JEL: G11, G13, G32

Jens Jackwerth is at the University of Konstanz, PO Box 134, 78457 Konstanz, Germany, Tel.: +49-(0)7531-88-2196, Email: jens.jackwerth@uni-konstanz.de. I thank Jim Hodder for helpful comments.

Delegated Portfolio Management and Risk-Taking of Hedge Funds

Introduction

Busy investors often delegate portfolio management to professional managers. They then might well ask if the risks assumed reflect the preferences of the manager more than their own. Most single-period models predict substantial variation in risk-taking across hedge fund value and time as managers react to incentive fees when exceeding high-water marks (Carpenter 2000, Basak, Pavlova, and Shapiro 2007, Kouwenberg and Ziemba (2007), and Buraschi, Kosowski, and Sritrakul 2014) or a lower dismissal boundary (Goetzmann, Ingersoll, and Ross 2003). Hodder and Jackwerth (2007) offer a model encompassing all these aspects plus endogenous shut-down in that the manager can close the fund and take up some outside option.[1]

So investors should simply measure risk-taking and check, by, say, working out the volatility of the hedge fund as a function of fund value and time. But as there are only 12 monthly returns in a year for any particular hedge fund, we have essentially two data points: risk-taking during the first half vs. the second half of the year. Clare and Motson (2009), Aragon and Nanda (2012), Ray (2012), and Li, Holland, and Kazemi (2019) find some evidence that hedge funds that perform poorly during the first half of a year take more risks during the second half. Buraschi, Kosowski, and Sritrakul (2014) also use the time-series of hedge fund returns but add a single-period structural model to investigate the impact of risk-taking on performance attribution (estimation of hedge fund alpha). Still, the empirical risk-taking surface remains shrouded in uncertainty.

[1] Hodder and Jackwerth (2011) even incorporate joint optimal management of the hedge fund and some correlated personal wealth account.

As an alternative, I estimate normalized risk-taking surfaces across many different funds, even when funds vary in investment opportunities, incentive contracts, preferences, and return distributions. Once returns are properly normalized by their time-series volatility, I can pool information cross-sectionally and precisely estimate normalized risk-taking surfaces. The intuition is much the same as in a standard t-test. Even though two variables (*x* and *y*) might have different distributions, once normalized, $x/\sigma(x)$ and $y/\sigma(y)$ both follow a t-distribution.[2]

Normalization works perfectly in the simple Merton (1969) model, where an investor (or a manager who receives a management fee only) optimally invests in a risky and a risk-free asset. When returns are divided by their time-series volatility, any risk-taking surface will be flat at 1.0, independent of the model parameters. Next, I simulate the richer Hodder and Jackwerth (2007) model of delegated portfolio management, while I vary parameters (expected returns, variances, incentive contracts, risk aversion) and modeling choices (power vs. exponential utility, normally vs. t-distributed returns). The normalized risk-taking surfaces no longer align perfectly across different funds but are still close to one another.

Thus, cross-sectional volatility of normalized returns appears to be a useful way to measure normalized risk-taking by hedge funds. Pooling a large cross-section of some 12,000 hedge funds allows non-parametric estimation of the normalized risk-taking surface. Empirically, that surface is almost flat, much more so than single-period models suggest. Expected confidence intervals are tight, at 0.99 and 1.01 for the 5^{th} and 95^{th} percentiles. Under most circumstances, hedge fund investors can take comfort in this empirical evidence that their fund managers do not wildly dial risk up or down in pursuit of their own gain.

---

[2] The mean of *x* and the mean of *y* need to be normally distributed.

The evidence indicates that only hedge funds operating far below their high-water marks take more risk around May and October. Such behavior might be described as "gambling for resurrection," in that it can catapult the manager away from a lower dismissal boundary if the gamble goes well. If not, the hedge fund investors (and not the manager) are saddled with the costs of the gamble. Interestingly, that increased risk-taking occurs somewhat before year-end (October instead of December), early enough so that the affected returns can be audited in time for the accounting year-end, when the fee income is determined. A similar but less definite event occurs in May, possibly for the smaller subset of hedge funds that determine their fee income on the basis of an accounting year-end in June.

Flat risk-taking is a notion contrary to most single-period models, but it is consistent with the multi-year version of Hodder and Jackwerth (2007). In that case, the manager decides on risk-taking when fee income is calculated at the end of each year. With such a longer-term perspective, dramatic shifts in risk-taking near a fund's high-water mark are suboptimal.[3] Note that gambling for resurrection, as I empirically find it, can still occur for low fund values in the multi-year model.

The cross-sectional estimation of normalized risk-taking surfaces is a universal tool for study of a variety of aspects of delegated portfolio management. Are mutual fund managers guided by

---

[3] Panageas and Westerfield (2009) study performance fee contracts (without management fees) with immediately resetting high-water marks in a continuous-time, infinite-horizon model. Immediately resetting high-water marks is not what occurs in practice; however, that approach does produce a risk-taking surface, that is flat and varies only when the fee income occurs at some finite horizon. Their model does not include gambling for resurrection as they do not model dismissal of the manager at low fund values.

tournament-style fund flow concerns, and take risks accordingly? Do prime brokers take risks on customer accounts so that they can manage their own inventory? Do pension funds take risks to insure their own long-term survival? In each case, the cross-sectional normalized risk-taking surface could provide empirical evidence to tell investors whether they need to be worried.

I describe methodology and data in Sections 1 and 2. Results are in Section 3. Section 4 concludes.

1. Methodology for normalized risk-taking surfaces

I introduce normalized risk-taking surfaces in three steps. In the Merton (1969) model, normalization works perfectly by design. In the richer Hodder and Jackwerth (2007) model, it works well in simulations for varying parameters. Finally, I implement the method empirically.

1.1. Normalization works perfectly in the Merton (1969) model

Merton (1969) studies an investor, but his approach works also for a fund manager whose fee income is a management fee only (and possibly some stake in the managed fund) but no performance fee. The manager can invest the fund value with proportion k in some risky, log-normally distributed asset and with proportion (1 - k) in the risk-free asset. The optimal proportion k is constant:

$$k = \frac{\mu - rf}{\gamma\, \sigma^2}, \tag{1}$$

where μ is the expected return on the risky asset, rf the risk-free rate, γ the risk-aversion of a power utility manager, and σ^2 the variance of the risky asset. The hedge fund return r_{HF} is:

$$r_{HF} = rf + k\,(r - rf), \tag{2}$$

where r is the realized return on the risky asset. I measure risk-taking as $\sigma_{HF}(r_{HF}) = k\, \sigma_r$, which is constant for a fund in the Merton framework, but could vary across hedge funds because of differing values for k. Next, I normalize returns $\left(r_{HFN} = \frac{r_{HF}}{\sigma_{HF}}\right)$ and calculate normalized risk-taking as the volatility of normalized returns, which will be 1.0 for each hedge fund at all fund values and times. In this simplified setting, funds can have differing risk-aversions γ, variances σ^2, and expected excess returns ($\mu - rf$), but normalization eliminates any effect of different k values (as long as σ_{HF} is estimated precisely). Pooling normalized hedge fund returns to compute cross-sectional normalized risk-taking works perfectly – in this setting normalized risk-taking is trivially 1.0 at all fund values and times.

1.2. The Hodder and Jackwerth (2007) model is a rich, single-period model of delegated portfolio management

In practice, we do not have these idealized conditions, but normalized risk-taking surfaces align across hedge funds with varying parameters even in a more realistic model. The Hodder and Jackwerth (2007) model adds considerable realism to the theoretical analysis of hedge fund risk-taking. It allows for a managerial stake in the hedge fund ($a = 0.10$ in the base case), a management fee ($b = 0.02$), and a performance fee ($c = 0.20$). Other base case parameters include a risky asset yielding 0.0778 per year with a variance of 0.05^2, and a risk-free asset returning 0.0578 per year. At each weekly time-step over a one-year period, the manager chooses the optimal proportion k to invest in the risky asset, where k varies from zero to five in steps of 0.01. The initial hedge fund value is 1.0, the same as the fund's high-water mark, and the lower dismissal boundary is 0.5. This dismissal boundary works like a maximum drawdown constraint; see Van Hemert, Ganz, Harvey, Rattray, Martin, and Yawitch (2020). If the fund value hits or breaches the lower boundary, the manager receives the pro-rated management fee earned until that point plus its own stake. All income is converted through a power utility function with risk aversion coefficient $\gamma = 4$.

There are 600 steps between the log of initial fund value and the log of the lower boundary, and twice that many above the log of initial fund value. The normal distributions of log returns are discretized into 121 values. To speed computation in this paper, I use 100 fund value steps and 21 discretization values without noticeable impact on results in the figures displayed below.

Figure 1 shows the resulting risk-taking surface, which displays the typical non-flat features of single-period models. Just below the high-water mark and toward the end of the year, the "option ridge" indicates high risk-taking as the manager tries to get the performance fee into the money. Once that goal is achieved, the manager greatly reduces risk-taking in an attempt to lock in the performance fee. At somewhat higher fund values, risk-taking ramps up to the Merton risk-taking level (see Eqn. 1) of 2.0 for the base case. Below the option ridge, there is also a "Merton flats" area, before a drop into the "valley of prudence," where the manager dials down risk-taking to avoid dismissal. Note the (tiny) "gambling for resurrection" just above the lower boundary near the end of the year. In that area, the manager takes more risk in hopes to move the fund back to higher fund values.

Fig 1. Risk-taking surface in the Hodder and Jackwerth (2007) model

Risk-taking surface for the base case (parameters in the text) of a hedge fund manager optimally investing into a risky and a risk-free investment over one year. Initial fund value and high-water mark are 1.0, and the manager is dismissed upon hitting the lower boundary of 0.5.

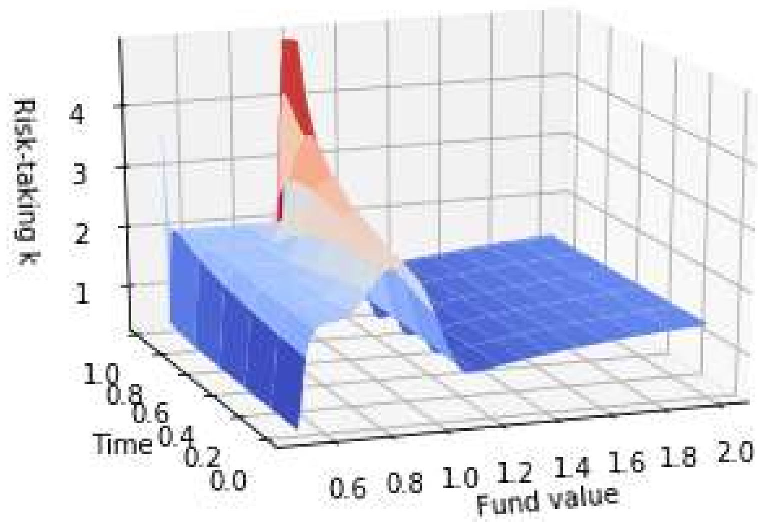

1.3. Normalization works well for simulations of the Hodder and Jackwerth (2007) model

In the Hodder and Jackwerth (2007) model, varying parameters leads to different risk-taking surfaces that are not easily comparable. To make comparisons possible, I introduce the normalized risk-taking surface, which looks very similar across hedge funds with varying parameters.

I normalize log returns at each grid point (identified by normalized log fund value relative to the high-water mark *X* and time *t*) by volatility across the whole surface σ_{HF}:

$$r_{HFN,X,t} = \frac{\log(1+r_{HF,X,t})}{\sigma_{HF}}. \quad (3)$$

In the model, I compute σ_{HF} in three steps. First, I gather nodal probabilities. The nodal probability is one at the initial log fund value $(X = 0)$ and time $(t = 0)$. I then move forward, one time-step at a time, use the return probabilities and the optimal risk-taking k at the grid point, and find the nodal probabilities at the next time-step. Nodal probabilities at each time-step sum to one. Second, I compute the nodal variance of (log of one plus return) = $k_{X,t}^2\, \sigma_r^2$ and take expectations for each time-step using the nodal probabilities. Finally, I take the average of the time-step variances and take the square-root to find σ_{HF}.

I measure local normalized risk-taking k^* as the volatility of normalized returns at that grid point:

$$k_{X,t}^* = \frac{k_{X,t}\, \sigma_r}{\sigma_{HF}}. \tag{4}$$

Figure 2 shows the normalized risk-taking surface for the base case across normalized fund values relative to the high-water mark and time. Centered around a normalized log fund value of log (1.0) = 0, it replicates the original risk-taking surface from Figure 1.

How similar are normalized risk-taking surfaces to the base case surface as we vary parameters in the Hodder and Jackwerth (2007) model? In Appendix A, I show surfaces as parameters are altered. Moderate parameter variations do not lead to drastic changes in normalized risk-taking. Stylized findings from the Hodder and Jackwerth (2007) model are clearly recognizable: option ridge, valley of prudence, ramp up to Merton flats, Merton flats, and even in many cases the small gambling for resurrection.

Fig 2. Normalized risk-taking surface in the Hodder and Jackwerth (2007) model

Normalized risk-taking surface for the base case (parameters in the text) of a hedge fund manager optimally investing into a risky and a risk-free investment over one year. The normalized log initial fund value and high-water mark are zero, and the manager is dismissed upon hitting the lower boundary of about -25 in normalized units of monthly return volatility.

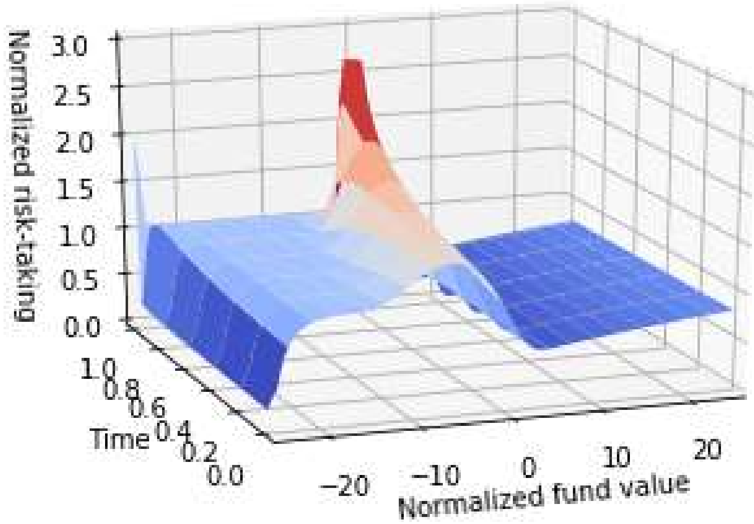

Rather than simply eyeball the normalized risk-taking surfaces in Appendix A, I measure expected absolute differences in risk taking *k* between the altered case and the base case as:

$$\Delta k = E[|k_{X,t} - k_{X,t} \text{ in the base case}|], \quad (5)$$

where I base the expectation on the base case probabilities of reaching a normalized grid point *(X,t)*. Further, I measure expected absolute differences in normalized risk-taking k^* as:

$$\Delta k^* = E[|k^*_{X,t} - k^*_{X,t} \text{ in the base case}|]. \quad (6)$$

In the Merton (1969) model with the parameters of the base case, $k_{X,t} = \frac{\mu - rf}{\gamma\sigma^2} = \frac{0.0778 \quad .0578}{4 \cdot 0.05^2} = 2.$

If, for example, one increases the expected excess return by 10% (or reduces the risk-aversion or the variance by 9.09%), that changes the risk-taking to 2.2. The expected absolute difference is Δk = 0.2 as the risk-taking surfaces of the Merton (1969) model are flat. The normalized risk-taking is always 1.0, however, so that the expected absolute difference in normalized risk-taking $\Delta k^* = 0$.

Table 1. Similarity of normalized risk-taking surfaces

Normalized risk-taking for the base case (Figure 2) is compared with the base case with one new parameter (Figures A2 through A11 in Appendix A). Differences in risk-taking Δk (expected absolute differences between k and k for the base case) and normalized risk-taking Δk^* (expected absolute differences between k^* and k^* for the base case) are recorded.

Figure	New parameter	Base case	E(abs(k - k_base))	E(abs(k^* - k^*_base))
A2	gamma $\gamma = 3$	$\gamma = 4$	0.6	0.1
A3	exponential utility	power utility	0.2	0.1
A4	variance = 0.07^2, i.e., Sharpe ratio = 0.29	variance = 0.05^2, i.e., Sharpe ratio = 0.40	0.8	0.1
A5	expected risky excess return = μ -rf = 0.027	μ -rf = 0.02	0.4	0.1
A6	t-distribution with 8 degrees of freedom	normal distribution	0.6	0.1
A7	own stake a = 0.16	a = 0.10	0.2	0.1
A8	management fee b = 0.11	b = 0.02	0.2	0.1
A9	performance fee c = 0.12	c = 0.20	0.2	0.1
A10	initial fund value w = 0.97	w = 1	0.3	0.1

Table 1 summarizes the parameter changes and reports differences in risk-taking between the base case and the altered cases. I vary the parameters in Table 1 in such a way that $\Delta k^* = 0.1$, so that I can compare changes in different parameters with one another. The management fee has an elasticity of only 0.02 (percent change in Δk^*/ percent change in the parameter), so its changes hardly affect normalized risk-taking surfaces. A second group of parameters (risk-aversion, variance, expected risky excess return, managerial stake, performance fee, and the choice of distribution) has elasticities of around absolute 0.2, so changes do not affect the normalized risk-taking surfaces much. Interestingly, changing the normal distribution to a fat-tailed t-distribution with 8 degrees of freedom is in this group.

One final parameter (initial fund value) has an elasticity of -3.33. This high value is caused by the presence of the option ridge, as it now matters if the hedge fund starts under water (initial fund value lower than the high-water mark). The high volatility area of the option ridge is then under-sampled, leading to lower normalized risk-taking. Luckily (foreshadowing the empirical results), option ridge does not seem to exist in the data. And for almost flat normalized risk-taking surfaces, the elasticity of initial fund value is close to zero. For the change from power utility to exponential utility I cannot compute an elasticity, but the small Δk^* of 0.1 suggests that considerable variation in preferences does not affect the normalized risk-taking surfaces much.

The purpose of this exercise is to demonstrate that normalized risk-taking surfaces are similar to each other for at least a moderate range of parameter adjustments. The results are from the single-period version of a theoretical model, but they do indicate that normalized risk-taking surfaces may be similar across funds with somewhat differing incentive fees or ownership stakes, for example. This is encouraging when one considers using cross-sectional data to empirically estimate such risk-taking surfaces. Furthermore, my empirical work uses hedge funds that are similar to each other in many dimensions (such as fee structure, style, currency). Hence, it is reasonable to expect that their normalized risk-taking surfaces are also fairly similar. Moreover, independent variation in the parameters is innocuous as I average across many (squared) normalized returns at each grid point *(X, t)* of the normalized risk-taking surface.

1.4. Empirical normalized risk-taking surfaces

The empirical work involves normalizing individual hedge fund returns and aggregating normalized risk-taking across hedge funds. For each hedge fund, I construct the cumulative returns that accrue to a dollar of initial fund value. I construct the high-water mark for the beginning of each year, which is the highest fund value reached so far, measured every three months (the so-called crystallization frequency). This high-water mark remains in place until the end of the year.[4]
I estimate each hedge fund's return volatility as the time-series volatility of logarithms of one plus monthly net-of-fee returns. For the normalization, I divide log returns by their time-series volatility to arrive at monthly normalized log returns. As I typically encounter some five years of hedge fund returns in the database, my volatility estimate is typically based on five paths across the (one-year) risk-taking surface. Similarly, I take logs of the ratio of fund values (that is, cumulative returns) and high-water marks, and normalize these values by the time-series volatility.

Next I aggregate risk-taking across similar hedge funds, whose normalized risk-taking surfaces should be fairly similar. Starting with the normalized returns at different normalized log fund values relative to the high-water marks *(X)* and time *(t)*, I compute the normalized risk-taking

---

[4] Elaut, Froemmel, and Sjoedin (2014) find a three-month crystallization frequency most common for hedge funds with the style of Commodity Trading Advisors. Using a 12-month crystallization frequency, which is commonly assumed in the literature, does not change results much at all. A more complicated alternative to compute the high-water mark is Buraschi, Kosowski, and Sritrakul (2014, p. 2848).

surface. For each month of the year, I collect all normalized returns. I specify X_s values from -10 through 10 in steps of 2, where I sample the surface (one unit corresponds to one standard deviation of monthly returns) and compute kernel-based weights with bandwidth h = 3, giving more importance to nearer returns and less to farther away returns. At $X_s = 0$, the fund value equals the high-water mark. Note that, in January, any return can start only at $X \leq 0$ as funds cannot start above the high-water mark.

The weights for each return i are probabilities as they are positive and sum to one:

$$p_i = NDF\left(\frac{X_i - X_s}{h}\right) / \sum_i NDF\left(\frac{X_i - X_s}{h}\right) \tag{7}$$

where *NDF* is the normal density function. Using the probabilities p_i, I compute the volatility of log returns around (X_s,t). I use similar weights but with a bandwidth of 0.82 to count the returns around (X_s,t).[5]

Normalized risk-taking seems to be unbiased. An upward bias in cross-sectional volatility might arise when the expected returns for individual hedge funds differ from the cross-sectional expected return at (X, t). A downward bias might arise if cross-sectional returns are positively correlated. Since I use normalized returns, the cross-sectional volatility should be 1.0 if all hedge funds share the same cross-sectional expected return and have uncorrelated returns. Empirically, the average cross-sectional volatility turns out to be 1.01, and the potential biases do not seem to matter much.

---

[5] The bandwidth governs how many neighboring returns are counted. The expected number of returns thus depends on the bandwidth. For a bandwidth of 0.82, the expected number of returns across the whole surface closely matches the total number of returns in the data.

2. Data

The empirical work uses a joint hedge fund database that includes Morningstar, Eurekahedge, BarclayHedge, HFR, TASS, CISDM, and Preqin. Names of managing companies are the basis for the merge, as in Hodder, Jackwerth, and Kolokolova (2014) and Joenvaara, Kauppila, Kosowski, and Tolonen (2020). I remove duplicates and different share classes of the same fund within each company by grouping the funds if their return correlations are above 0.99. Within each group, I keep the fund with the longest time-series of returns. The data consist of 5.5 million monthly returns of 78,880 funds (for 11,592 management companies) from January 1986 through May 2018.

Standard filters for the data involve a requirement of consecutive returns during at least 24 months. I use only funds that report in U.S. dollars (reducing the sample by 7%). As the great majority of managing companies offer multiple share classes in the same strategy, this picks the USD share class, which is also the most common. To address survivorship bias, I keep both dead and live hedge funds, which is possible for funds starting around 1993. To address backfill bias, I remove the first 12 returns to mitigate back-fill bias; see Kosowski, Naik, and Teo (2007) and Teo (2009). The resulting data cover January 1994 through May 2018.

Further filters ensure that hedge funds are reasonably homogeneous and share similar normalized risk-taking surfaces. Hedge funds within the same style should be more homogeneous, and I filter for equity and equity long/short funds (reducing the sample by 19%). Using all styles simultaneously yields a very similar risk-taking surface. The management fee needs to be between 1% and 2% (reducing the sample by 16%). Allowing only a management fee of 1.5% yields a very similar risk-taking surface. The performance fee needs to be 20% (reducing the sample by 33%), and only hedge funds with a high-water mark enter the sample (reducing the sample by 9%).

The filtering shrinks the sample considerably, from 5.5 million returns to 0.7 million, as we see in Table 2, Panels A and B. Minimum and maximum monthly returns are determined by the decision to exclude returns below -100% and above 200%. Mean and standard deviation are in line with other hedge fund studies, at around 0.64% and 4.69%. Skewness and excess kurtosis move closer to zero in the process of filtering, and even more so in Panel C for the filtered normalized returns with skewness of -0.19 and excess kurtosis of 5.36. Then, the standard deviation is close to one (1.0126), and normalized returns do not seem substantially biased.

Table 2. Descriptive statistics

Descriptive statistics for monthly returns are number of observations (Obs.), minimum, maximum, mean, standard deviation (STD), standardized skewness, and excess standardized kurtosis. Panel A shows the statistics for all returns, Panel B for the filtered returns, and Panel C for the filtered normalized returns.

Panel A. All monthly returns

Obs.	Min	Max	Mean	STD	Skewness	Kurtosis
5480764	-1.00	1.86	0.0064	0.0469	1.28	54.28

Panel B. Filtered monthly returns

Obs.	Min	Max	Mean	STD	Skewness	Kurtosis
720966	-0.87	1.72	0.0074	0.0539	0.71	25.86

Panel C. Filtered monthly normalized returns

Obs.	Min	Max	Mean	STD	Skewness	Kurtosis
720966	-12.10	14.79	0.1741	1.0126	-0.19	5.36

3. Results

The primary result is the empirical normalized risk-taking surface shown in Figure 3, where I set areas to zero with fewer than 1,000 returns around a grid point (X_s, t). Normalized risk-taking empirically bears no resemblance to standard single-period models such as Hodder and Jackwerth

(2007) in Figure 2.[6] The option ridge near the high-water mark does not exist empirically, and the associated ramp-up to Merton flats for higher fund values also does not appear. This puts the empirical studies of Aragon and Nanda (2012), Ray (2012), Buraschi, Kosowski, and Sritrakul (2014), and Li, Holland, and Kazemi (2019) into perspective. That is, if the effects they are looking for (increased risk-taking after initial losses for three papers, and an option ridge for the other) do not exist in the data, then the empirical evidence will be weak.

Figure 3. Empirical normalized risk-taking

Normalized risk-taking (Z-axis) across normalized fund values relative to high-water marks (X-axis) and time (Y-axis).

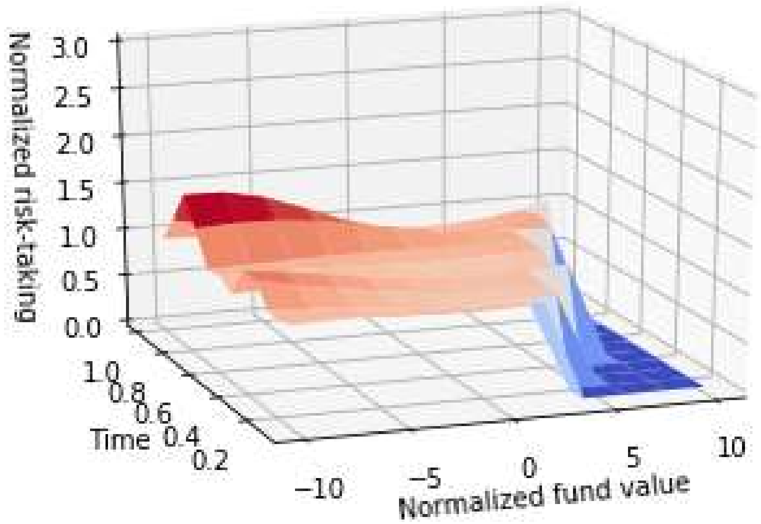

[6] The result does not change when I base volatilities on log excess returns instead of log returns.

Instead of a Merton flats area below the high-water mark, there appear ridges of increased normalized risk-taking for low fund values. The October ridge is more pronounced than the May ridge.

Note that the risk-taking surface applies for both recessions and boom periods. During recessions, more paths pass through low fund values, and during booms, more paths pass through high fund values, but all paths occur on the same risk-taking surface.

3.1. Confidence intervals are tight

For bootstrap confidence intervals I follow Haerdle (1990, pp. 103), and bootstrap 100 versions of my database with replacement from the universe of hedge funds (12,060 unique hedge funds) so that I keep intact possible time-series patterns within hedge funds. At each grid point I sample the 5th and 95th percentiles of the estimated normalized risk-taking surface. The results are undistinguishable from Figure 3. I measure the expected absolute differences from the base surface, and find $\Delta k^* = 0.01$ (see Eqn. 6) both for the lower and the upper percentile. The confidence interval is exceedingly tight (the surface itself is almost flat at 1.0) due to the large number of pooled returns (721,000).

The quality of the surface depends on the number of cross-sectional normalized returns that I can use. Figure 4 shows this count, which is concentrated in January just below $X = 0$ with some 30,000 returns and then spreads out as time passes. By December, most hedge funds show fund values that are two monthly standard deviations up from the initial fund value. Even for very low fund values, I still collect more than 1,000 returns at each grid point. For high fund values at the beginning of the year, very few returns are recorded. Because the high-water mark resets, there cannot be any January observations greater than zero for a normalized log fund value divided by its high-water mark (i.e., $X > 0$ is not possible at the beginning of January). Risk-taking surfaces (and also the count surface in Figure 4) are set to zero whenever there are fewer than 1,000 returns in the cross-section.

Figure 4. Number of normalized returns

Number of normalized returns (Z-axis) across normalized fund values relative to high-water marks (X-axis) and time (Y-axis).

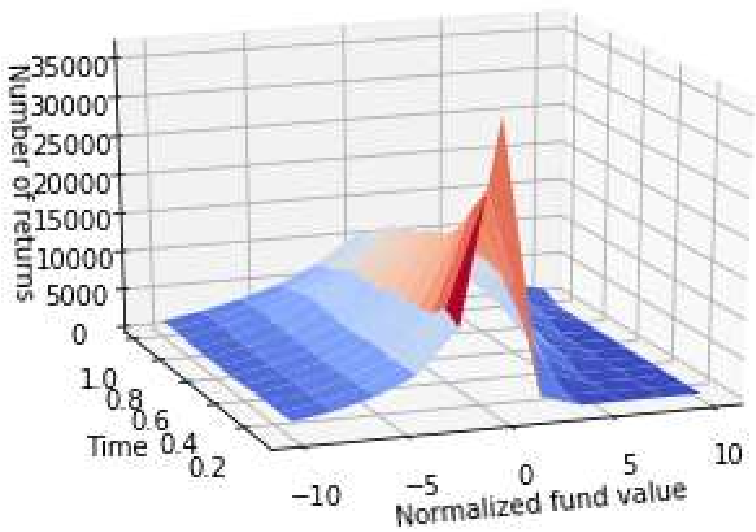

3.2. Empirical results are consistent with multi-year managerial incentives

The missing option ridge is likely due to the (perhaps implicit) multi-year nature of incentives for hedge fund managers. Managers do not optimize risk-taking solely with end-of-year compensation in mind. They also take into account the expected value of future years' compensation. Hodder and Jackwerth (2007) cover such multi-year horizons, when the resulting risk-taking surfaces do not show an option ridge associated with the high-water mark. In a related vein, Panageas and Westerfield (2009) point out that an infinitely far away performance fee with immediately resetting high-water marks leads to an initially flat risk-taking surface. Their model is not particularly realistic, as hedge fund compensation is a sequence of annual payoffs, as modeled in Hodder and Jackwerth (2007), rather than a single far-away payoff.

Elevated risk-taking at low fund values is consistent with Hodder and Jackwerth (2007). Such behavior could be due to new opportunities arising once the manager has been dismissed at the lower boundary or decides itself to quit (perhaps to start a new hedge fund with a reset high-water mark, or to enter a new work contract, or to enjoy leisure). Interestingly, the increased risk-taking is concentrated in October with a secondary peak in May. This timing could be related to the audit, which, by best practice, should provide the basis for determining manager compensation (Bruebaker, 2009, p. 53).

Table 3. Audit dates

Months of audit dates reported in the TASS hedge fund database.

Month	Observations	Proportion
Jan	124	0.02
Feb	95	0.01
Mar	197	0.03
Apr	100	0.01
May	60	0.01
Jun	445	0.06
Jul	62	0.01
Aug	108	0.02
Sep	367	0.05
Oct	78	0.01
Nov	92	0.01
Dec	4352	0.62
not reported	971	0.14
Total	7051	1.00

Table 3 reports audit months (audit date is typically the last day of the month) from the TASS database, which is my only database that provides this information. 62% of all hedge funds have audit dates in December. Another 6% are in June, and 14% do not report an audit date. Accounting research documents a preference for year-end audits and discounts for audits at other times of the year; see Ng, Tronnes, and Wong (2017). As the returns affected by changed risk-taking need to be included in the audit to affect compensation, managers would seek to take risks somewhat before the June and December audit dates, consistent with my empirical finding of increased risk-taking in May and October.[7]

3.3. The multi-period Hodder and Jackwerth (2007) model is consistent with the empirical normalized risk-taking surface

Figure 5 shows the normalized risk-taking surface for the first year of a five-year contract in the Hodder and Jackwerth (2007) model. All parameters are as in the base case, but now I model five years of weekly risk-taking decisions. The manager receives at each year-end the management and the performance fee, (which is based on the high-water mark). Additionally, at the end of the fifth year, the manager's own stake is cashed out. The lower boundary works as in the single-period model.

The option ridge is not visible for fund values just below the high-water mark, and there is virtually no ramp-up to Merton flats for higher fund values. There is a pronounced gamble for resurrection along the lower boundary. Given the five-year perspective, it is worthwhile to take more risks when fund values are low, in the hope of moving fund values back to the high-water mark. Even closer

---

[7] An alternative explanation for the May and October elevated risk-taking does not seem to be likely. If high risk-taking hedge funds in May and October leave the sample in June and November, then risk-taking might be lower at that time. Yet, when I filter out intra-yearly leavers and use only hedge funds with a full calendar year of returns, the risk-taking surface looks just the same as before, including the May and October elevated risk-taking.

to the lower boundary, the valley of prudence opens up again, which is not visible empirically in Figure 3. This might be attributable to voluntary reporting of hedge funds in the databases, as poorly performing hedge funds often stop reporting their returns. Also, the lower boundary in reality is not a hard boundary but rather a wider region of the increased likelihood of severe fund outflow in reaction to poor hedge fund performance, which the hedge fund might counter with lower risk-taking.

Fig 5. Multi-year normalized risk-taking surface in the Hodder and Jackwerth (2007) model

Normalized risk-taking surface for the base case (parameters in the text) of a hedge fund manager optimally investing into a risky and a risk-free investment over the course of five years. The normalized initial fund value and high-water mark are zero, and the manager is dismissed upon hitting the lower boundary of about -30 in normalized units.

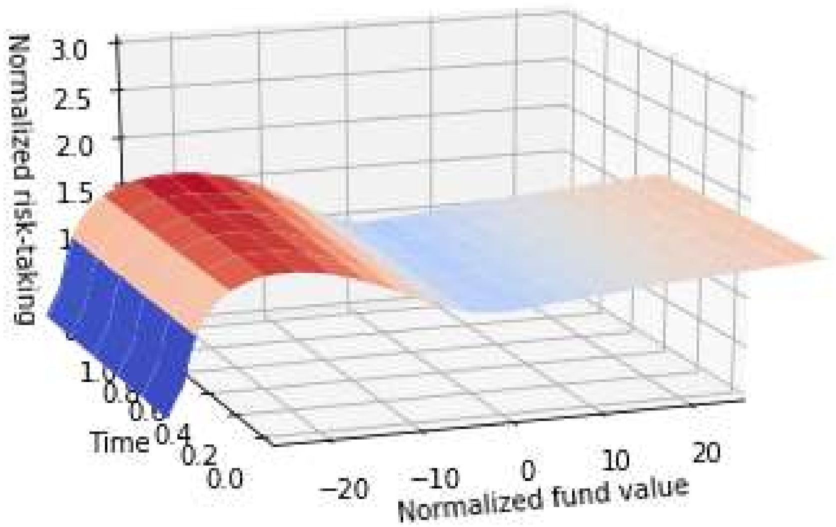

Empirically, we find some gambling for resurrection at low fund values in May and October. The fine structure of such increased risk-taking is absent from the Hodder and Jackwerth (2007) model, which does not consider proximity to the audit date.

3.4. Empirical normalized risk-taking surfaces are very robust

Results do not change much when I change the methodology. I checked by using excess returns instead of returns; by expanding the sample and allowing for all hedge fund styles (as opposed to only equity and equity long/short in the base case); by changing the crystallization frequency from three months to 12 months and one month; and by using only hedge funds with 1.5% management fee (1% through 2% in the base case).

As I filter more strictly, I lose around three-quarters of my return sample. To still see as much of the resulting normalized risk-taking surfaces as in Figure 3 for the base case, I cut the fringes when there are fewer than 100 returns (instead of 1,000 in the base case) around *(X, t)*.

Figure 6. Empirical risk-taking – initial fund value is same as the high-water mark

Normalized risk-taking (Z-axis) across normalized fund values divided relative to high-water marks (X-axis) and time (Y-axis).

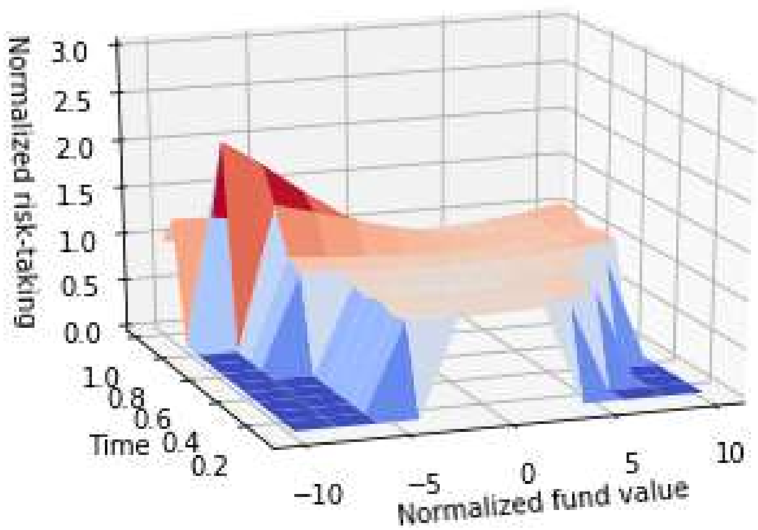

Figure 6 shows the normalized risk-taking surface if I use only hedge funds that start at the high-water mark (254,000 returns instead of 721,000 in the base case[8]). The surface is similar to the base case in Figure 3 except for the missing returns below the high-water mark near the start of the year. Figure B1 in Appendix B shows the same surface for hedge funds with a management fee of exactly 0.015 (250,000 returns). The surface is indistinguishable from that in Figure 3. Figure B2 shows the surface for large hedge funds (assets under management above the median value of USD 56.7

[8] This 65% reduction is in line with a simple four-step binomial tree (matching the base case crystallization frequency of three months) with equal probabilities, where 62% of returns are lost.

million; 285,000 returns) and Figure B3 for small hedge funds (assets under management below the median value; 235,000 returns). The surfaces are very similar to the one in Figure 3.

4. Conclusion

As wealthy investors delegate their portfolio management to professional managers, it is reasonable for them to be concerned about whether the risks the funds take possibly reflect the preferences of the money manager more than their own. Single-period models suggest managerial risk-taking at hedge funds varies widely across fund values and time. Investors could try to measure such risk-taking (say, in terms of volatility) at a fund, but this is hard as the time-series of fund returns consists of only a limited number of monthly returns that do not trace out the complete risk-taking surface.

I show that normalized risk-taking surfaces align well across different funds, even when hedge funds vary in investment opportunities, assumed preferences, or return distributions. Such normalization thus allows me to pool information cross-sectionally and precisely estimate normalized risk-taking surfaces.

These turn out to be much flatter than single-period models suggest. There is no empirical evidence of increased risk-taking just below the high-water mark or of attempts at somewhat higher values to lock in performance fees by reducing risk-taking. The multi-year Hodder and Jackwerth (2007) model can account for such flat risk-taking around the high-water mark, and the multi-year nature of managerial incentives seems to explain the empirical results.

There is, however, evidence of more risk-taking at very low fund values, in line with theoretical predictions of gambling for resurrection. The timing in May and October is just before audit dates in June and December on which compensation is based.

The cross-sectional estimation of normalized risk-taking surfaces is a basic tool that allows study of many aspects of delegated portfolio management. Are mutual fund managers really guided by tournament style fund flow concerns, and take risks accordingly? Do prime brokers take risks on customer accounts as they manage their own inventory? Do pension funds take risks to insure their own long-term survival? In all these cases, time series for each fund are short, which makes it hard to study risk-taking. Yet the cross-sectional normalized risk-taking surface, which pools information across many funds, can empirically indicate whether investors need to worry.

References

Aragon, George O., and Vikram Nanda, 2012, On Tournament Behavior in Hedge Funds: High Water Marks, Fund Liquidation, and Managerial Stake, *Review of Financial Studies* 25, 937–974.

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

Bruebaker, Gary, 2009, Principles and Best Practices for Hedge Fund Investors, *Report of the Investors' Committee to the President's Working Group on Financial Markets*, 1–68.

Buraschi, Andrea, Robert Kosowski, and Worrawat Sritrakul, 2014, Incentives and Endogenous Risk Taking: A Structural View on Hedge Fund Alphas, *Journal of Finance* 69, 2819–2870.

Carpenter, Jennifer N., 2000, Does Option Compensation Increase Managerial Risk Appetite? *Journal of Finance* 55, 2311–2331.

Clare, Andrew D., and Nick E. Motson, 2009, Locking in the Profits or Putting it All on Black? An Empirical Investigation Into the Risk-Taking Behaviour of Hedge Fund Managers, *Journal of Alternative Investments* 12, 7–25.

Elaut, Gert, Michael Froemmel, and John Sjoedin, 2014, Crystallization – the Hidden Dimension of Hedge Funds' Fee Structure, Working paper, Ghent University.

Goetzmann, William N., Jonathan E. Ingersoll, and Stephen A. Ross, 2003, High-Water Marks and Hedge Fund Management Contracts, *Journal of Finance* 58, 1685–1717.

Haerdle, Wolfgang, 1990, *Applied Nonparametric Regression* (Cambridge University Press, Cambridge, UK).

Hodder, James E, and Jens Carsten Jackwerth, 2007, Incentive Contracts and Hedge Fund Management, *Journal of Financial and Quantitative Analysis* 24, 811–826.

Hodder, James E., and Jens Carsten Jackwerth, 2011, Managerial Responses to Incentives: Control of Firm Risk, Derivative Pricing Implications, and Outside Wealth Management, *Journal of*

Banking and Finance 35, 1507–1518.

Hodder, James E., Jens Carsten Jackwerth, and Olga Kolokolova, 2014, Recovering Delisting Returns of Hedge Funds, *Journal of Financial and Quantitative Analysis* 49, 797–815.

Joenvaeaerae, Juha, Mikko Kauppila, Robert Kosowski, and Pekka Tolonen, 2020, Hedge Fund Performance: Are Stylized Facts Sensitive to Which Database One Uses? *Critical Finance Review* forthcoming.

Kosowski, Robert, Narayan Y. Naik, and Melvyn Teo, 2007, Do Hedge Funds Deliver Alpha? A Bayesian and Bootstrap Analysis, *Journal of Financial Economics* 84, 229–264.

Kouwenberg, Roy, and William T. Ziemba, 2007, Incentives and Risk Taking in Hedge Funds, *Journal of Banking and Finance* 31, 3291–3310.

Li, Ying, A. Steven Holland, and Hossein B. Kazemi, 2019, Duration of Poor Performance and Risk Shifting by Hedge Fund Managers, *Global Finance Journal* 40, 35–47.

Merton, Robert C., 1969, Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case, *Review of Economics and Statistics* 51, 247–257.

Ng, Hooi Ying, Per Tronnes, and Leon Wong, 2017, Audit Seasonality and Pricing of Audit Services: Theory and Evidence from a Meta-Analysis, *Journal of Accounting Literature* 40, 16–28.

Panageas, Stavros, and Mark M. Westerfield, 2009, High-Water Marks: High Risk Appetites? Convex Compensation, Long Horizons, and Portfolio Choice, *Journal of Finance* 64, 1–36.

Ray, Sugata, 2012, The Downside of High Water Marks: An Empirical Study, *Journal of Investment Management* 10, w/o page numbers.

Teo, Melvyn, 2009, The Geography of Hedge Funds, *Review of Financial Studies* 22, 3531–3561.

Van Hemert, Otto, Mark Ganz, Campbell R. Harvey, Sandy Rattray, Eva Sanchez Martin, and Darrel Yawitch, 2020, Drawdowns, Working paper, Duke University.

Appendix A

Figure A1. Normalized risk-taking, base case

Normalized risk-taking in the Hodder and Jackwerth (2007) model across normalized fund value and time. Parameters are time t = 1, number of steps n = 52, size of the normal approximation nn = 1+2·10, number of steps between HWM and lower barrier x = 100, initial fund value w = 1, lower boundary LB = 0.5, HWM = 1, rf = 0.0578, μ = 0.0778, variance = 0.05^2, risk aversion γ = 4, own stake a = 0.10, management fee b = 0.02, and performance fee c = 0.20. This figure is identical with Figure 2 in the main text.

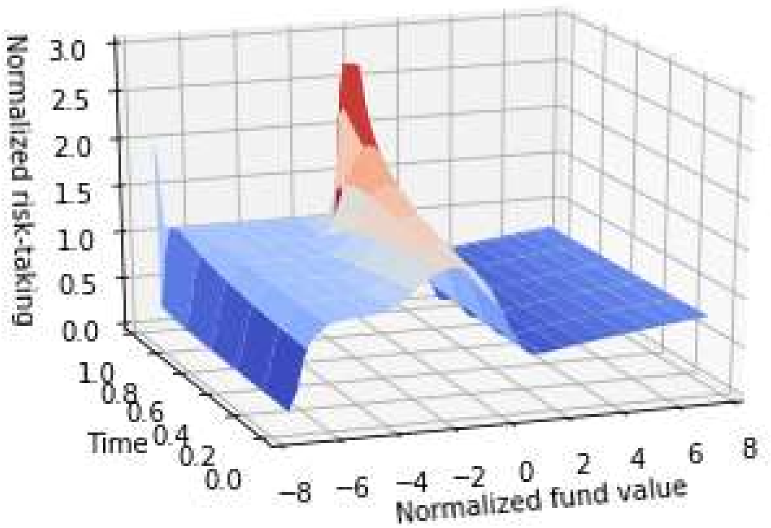

Figure A2. Normalized risk-taking, risk aversion γ = 3 instead of 4

Normalized risk-taking in the Hodder and Jackwerth (2007) model across normalized fund value and time. Parameters are as in the base case in Figure A1 except for risk aversion.

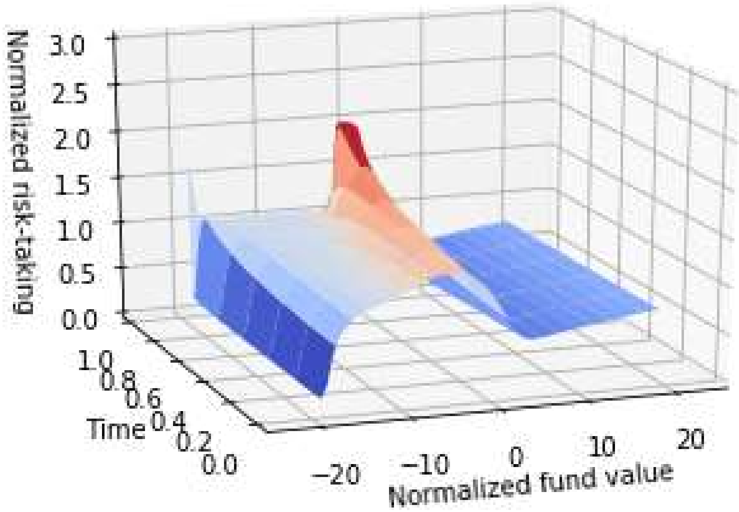

Figure A3. Normalized risk-taking, using exponential utility instead of power utility

Normalized risk-taking in the Hodder and Jackwerth (2007) model across normalized fund value and time. Parameters are as in the base case in Figure A1 except for an exponential utility function. Absolute risk aversion is 4/0.12 to roughly match the relative risk aversion of the base case.

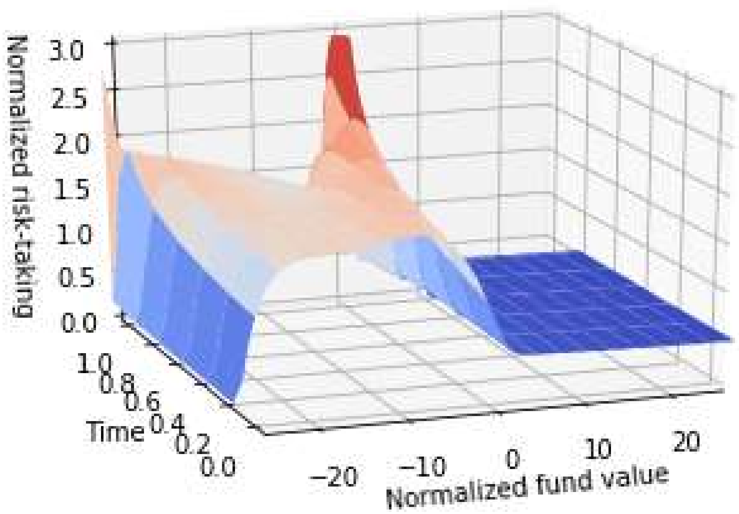

Figure A4. Normalized risk-taking, variance = 0.07^2 instead of 0.05^2 (i.e., Sharpe ratio = 0.29 instead of 0.40)

Normalized risk-taking in the Hodder and Jackwerth (2007) model across normalized fund value and time. Parameters are as in the base case in Figure A1 except for variance (Sharpe ratio).

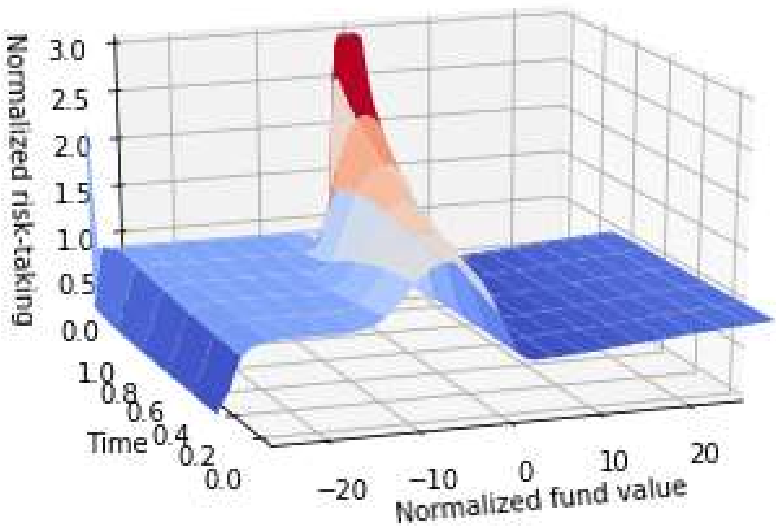

Figure A5. Normalized risk-taking, expected risky return μ = 0.0848 instead of 0.0778

Normalized risk-taking in the Hodder and Jackwerth (2007) model across normalized fund value and time. Parameters are as in the base case in Figure A1 except for expected risky return μ.

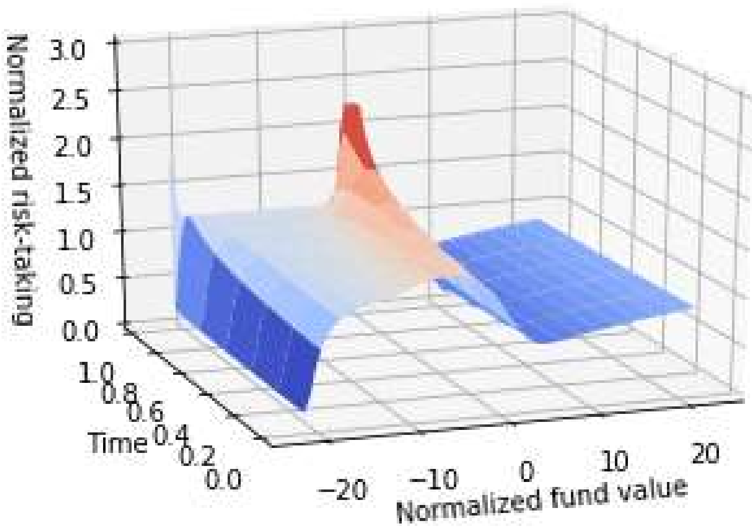

Figure A6. Normalized risk-taking, t-distribution with 8 degrees of freedom instead of normal distribution (corresponds to degrees of freedom tending to infinity)

Normalized risk-taking in the Hodder and Jackwerth (2007) model across normalized fund value and time. Parameters are as in the base case in Figure A1 except for a t-distribution.

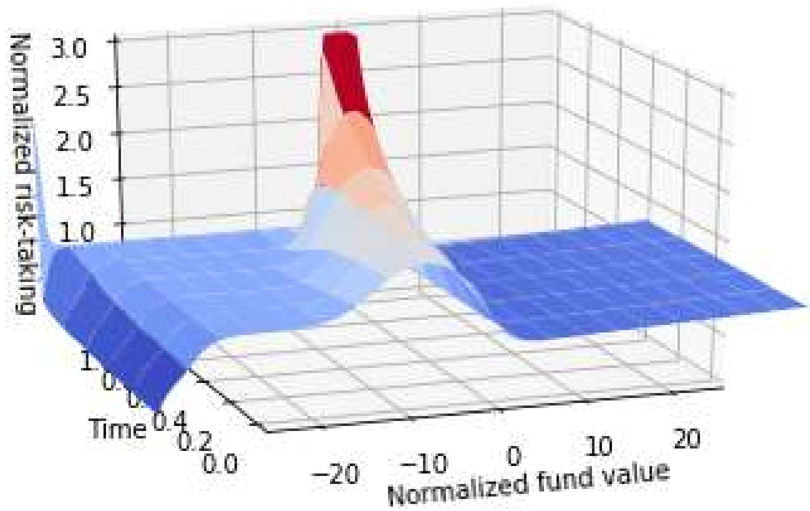

Figure A7. Normalized risk-taking, own stake a = 0.16 instead of 0.10

Normalized risk-taking in the Hodder and Jackwerth (2007) model across normalized fund value and time. Parameters are as in the base case in Figure A1 except for own stake a.

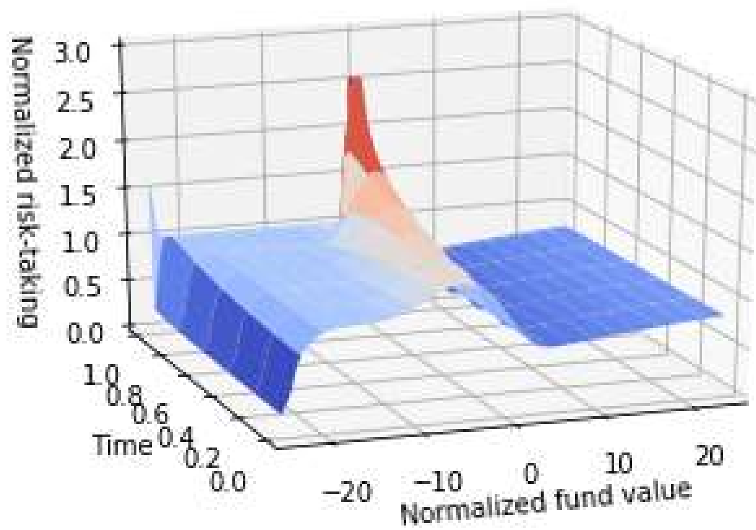

Figure A8. Normalized risk-taking, management fee b = 0.11 instead of 0.02

Normalized risk-taking in the Hodder and Jackwerth (2007) model across normalized fund value and time. Parameters are as in the base case in Figure A1 except for management fee b.

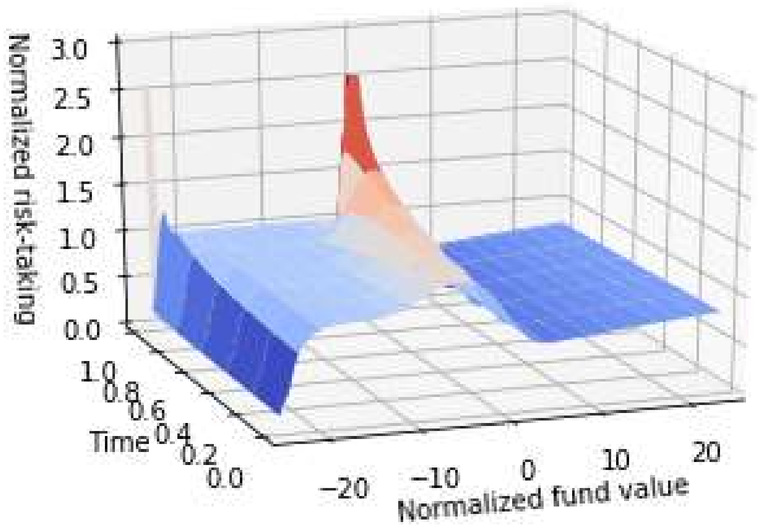

Figure A9. Normalized risk-taking, performance fee c = 0.12 instead of 0.20

Normalized risk-taking in the Hodder and Jackwerth (2007) model across normalized fund value and time. Parameters are as in the base case in Figure A1 except for performance fee c.

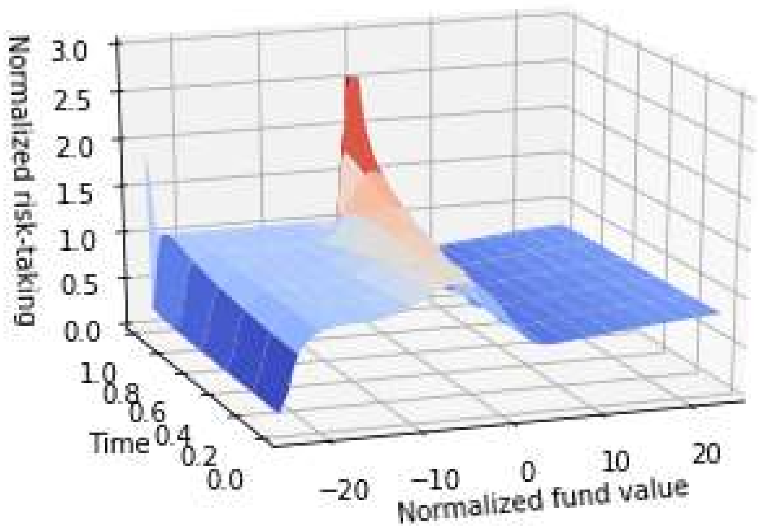

Figure A11. Normalized risk-taking, initial fund value w = 0.97 instead of 1

Normalized risk-taking in the Hodder and Jackwerth (2007) model across normalized fund value and time. Parameters are as in the base case in Figure A1 except for initial fund value w.

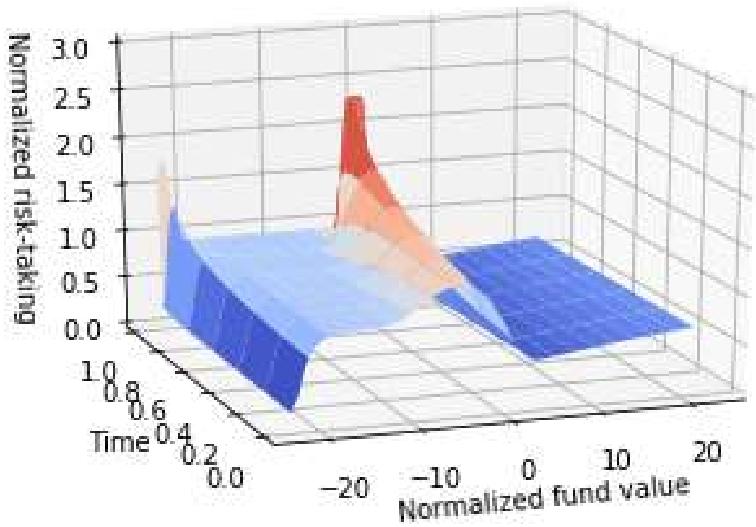

Appendix B

Figure B1. Empirical risk-taking, only hedge funds with management fee = 0.015

Normalized risk-taking (Z-axis) across normalized fund values divided relative to high-water marks (X-axis) and time (Y-axis).

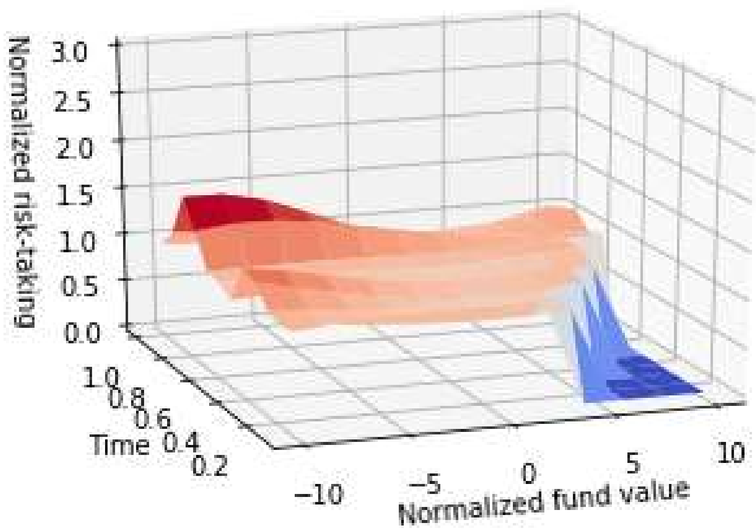

Figure B2. Empirical risk-taking, only large hedge funds

Normalized risk-taking (Z-axis) across normalized fund values divided relative to high-water marks (X-axis) and time (Y-axis).

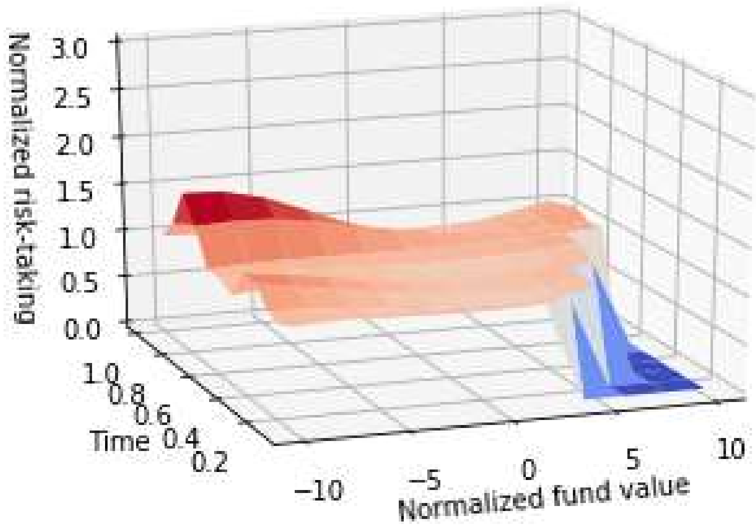

Figure B3. Empirical risk-taking, only small hedge funds

Normalized risk-taking (Z-axis) across normalized fund values divided relative to high-water marks (X-axis) and time (Y-axis).