Burgiss Applied Research: Working Paper

March 2018

Budgeting for Capital Calls: A VaR-Inspired Approach

Luis O'Shea and Vishv Jeet

In this paper we discuss the importance, and challenges, of forecasting private capital cash flows. Particularly important is forecasting capital calls, since they constitute liabilities for the investor. While it is useful to know the expected capital calls arising from an investor's portfolio, it is more important to estimate a likely upper bound on those calls, since this will determine the reserves needed in order to safely service the calls. To this end, we introduce a new concept, namely that of *maximum probable contributions* — a statistic which, subject to a user-specified confidence level, serves as such an upper bound. We explore in detail a historical methodology for its computation, illustrate typical model predictions, and document its out-of-sample performance during backtesting on both funds and portfolios of funds.

1 Introduction

Investors with commitments to private capital funds need to maintain sufficient reserves to service capital calls arising from those funds. If their reserves are insufficient, they may be unable to meet those capital calls (or perhaps, will be forced to sell illiquid assets at suboptimal prices). Conversely, excessive reserves will reduce the amount of capital in the ground, and consequently will result in the investor forgoing some of the expected returns arising from private capital. This paper proposes a systematic way to strike a balance between the two extremes of insufficient reserves and excess caution.

Let us suppose an investor considers these questions with respect to a fixed horizon, such as one quarter: How much capital should one have on hand to service capital calls in the next quarter? A possible answer to this question is that one should keep on hand the *expected*[1] capital calls (henceforth, contributions) in the next quarter. A moment's thought, however, shows that this is insufficient. Consider a simple example: suppose it is equally likely that in the next quarter contributions will be either \$2M or zero, then the expected contributions are \$1M. Maintaining reserves of \$1M means the probability that the investor will be unable to service their capital calls is 50%, a risk which is not likely to be acceptable. On the other hand, if one demands that the risk of being unable to meet one's capital calls be precisely zero, then the solution is simple: ensure reserves are at least equal to uncalled capital. This absolute certainty comes at a steep cost, namely that of diluting the expected returns of one's investments (since so much capital needs to kept in a liquid form). Informally it seems that a compromise is needed, whereby one keeps reserves at such a level that capital calls can usually be met, while allowing for the possibility that occasionally additional funds will be needed.[2] This idea, stated more formally, is the central concept discussed in this paper: given a level of confidence chosen by the investor (say, 95%), and a horizon (say, a quarter), we seek to estimate the amount of reserves necessary so that with probability 95% they will cover the capital calls in the next quarter. We call this amount the *95% maximum probable contribution (MPC)* of the investor's portfolio.

Our goal in this paper is to propose a methodology for computing MPC and to explore how the measure performs. In section 2 we give an introduction to the topic of probabilistic bounds (such as MPC and VaR) and in section 3 we list the (essentially non-existent) literature related to this idea with regard to private capital. In section 4 we describe, in detail, our methodology for estimating MPC, in section 5 we illustrate typical model predictions for various fund types and ages, in section 6 we backtest our model, and in section 7 we present our conclusions. Appendix B describes the data used for our model predictions and backtesting and appendix A provides more background on our backtesting methodology.

2 Risk and Probabilistic Bounds on Contributions

According to industry legend, around 1990 the chairman of JP Morgan requested that every day the risk of the bank's trading portfolio be summarized in the form of a single number included in a report delivered to his desk by 4:15pm. That number was the one-day 95% value-at-risk (VaR). The idea behind this number is that while a portfolio's value may be known today, at some future point in time (such as, tomorrow) its value can be best be described via a probability density function (PDF). Riskier portfolios would have more disperse PDFs than less risky ones. Thus the standard deviation of the PDF would indicate how risky the portfolio is. Indeed, if market returns were normally-distributed and all assets were linear, then the standard deviation would suffice. However neither of these assumptions are correct (for example, options have payoffs which are non-linear). These facts introduce the possibility of losses that exceed what one would anticipate based on standard deviations alone. For example it is possible for two portfolios with equal standard deviations to have very different probabilities of large losses. What JP Morgan's chairman needed was a measure of *tail risk*. Thus the concept of VaR was born — a quantile from the predicted PDF of profits and losses. For example, on

[1] We are using the term *expected* in its statistical sense. For example, if one flips a coin (once) then the expected number of heads is 0.5 (even though the observed number will always be either zero or one, and never one half!) One way of thinking of this is that it is the expected value is the long-run average. Equivalently, the expected value is the average over many identical realizations (coin flips, in the previous case). So the "expected capital calls" (over some period), are the average capital calls (over that period) from a large set of essentially identical funds.

[2] This compromise could be further rationalized as follows. Perhaps reserves are particularly liquid investments. If those liquid reserves prove insufficient then less liquid assets can be sold, perhaps at somewhat below market price. While evidently undesirable, this may be acceptable if it happens sufficiently rarely.

19 out of 20 days (95% of the time) JP Morgan's trading losses should be less than the reported one-day 95% VaR number.

Turning to private capital, investors face problems that are similar to those faced by JP Morgan (except related to cash flows rather than returns): the investor has a set of commitments which over the next quarter (say) will generate some set of capital calls. These calls cannot be known with certainty, thus the most one can hope for is a PDF of possible contributions (see figure 1 for an example of what such a PDF might look like). Note that such a distribution is very far from normal since contributions can only be positive, a contribution of precisely zero in any quarter is fairly common (i.e., the PDF is "zero-inflated"), and it decays to zero slower than a normal distribution.[3] In the PDF of figure 1 one can see a large probability of contributions being zero over the next period (this is the tall bar on the left of the histogram, and happens about 30% of the time in this example) as well as the exponential decay in probabilities for larger contributions. The expected (or mean) contribution is marked in blue. As can be seen, if one were to maintain reserves equal to that number then in about 50% of periods the reserves would be insufficient! Also shown (in red) is the 95th percentile (which is what we called the 95% MPC, above); by definition maintaining reserves at that level will result in them being sufficient on 95% of all periods. However note that the 95% MPC is about twice the expected contribution; this illustrates the wide discrepancy between these two numbers. An investor could, of course, try to set the level of their reserves based on a combination of expected contributions and a rule of thumb; for example, one could set reserves at twice the expected contributions. However, as we shall see in later sections, even this is likely to serve the investor poorly, on account of the fact that MPCs diversify across funds; thus depending on the sizes and ages of the funds in a portfolio the multiple of expected contributions that one should employ will vary — perhaps doubling the expected cash flows will be appropriate for certain portfolios at certain points of time, but in other situations it could be significantly different.

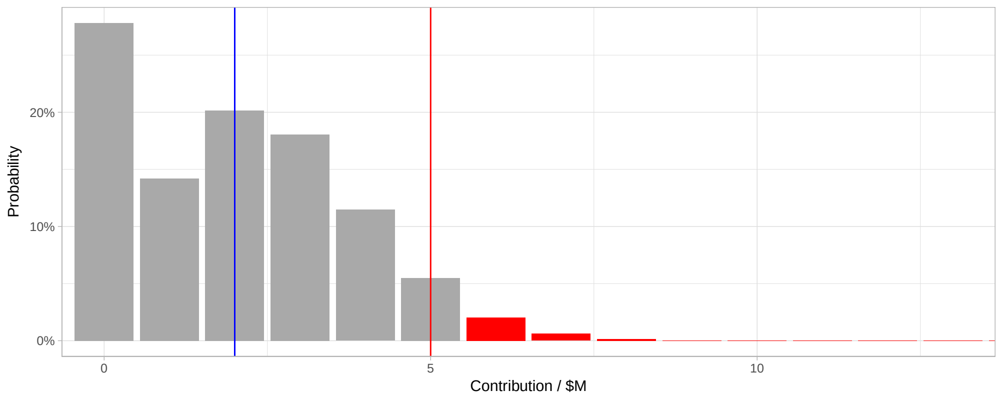

Figure 1: PDF of contributions to a private capital portfolio
The above figure is a histogram of possible contributions to a private capital portfolio. The blue line represents the expected contribution, the red line represents the 95% MPC contribution, and the red bars are contributions which exceed the MPC.

The above discussion should make clear that an estimate of MPC could be very useful for managing a portfolio of private capital funds. Furthermore computing such a measure is simple, assuming one has access to the PDF of possible contributions in the next period (as illustrated in figure 1). Thus estimating MPC comes down to predicting the PDF of future contributions. Much like when estimating VaR, there are two broad methodologies for predicting such a PDF: *historical* and *model-based*. Briefly, the distinction between these two approaches is as follows. In both VaR and MPC one must model how the market is likely to behave (VaR is concerned with returns, MPC with cash flows). One can impose a parametric form on this behavior, in which case one must estimate the parameters of this form; this is a model-based methodology. Alternatively, one can simply use historical data directly as a model for how the market behaves; this is a historical methodology. In this document we will focus almost exclusively on a historical methodology.

[3] Although it eventually goes to zero since contributions are capped by uncalled capital.

3 Literature Review

To the best of our knowledge MPC-like upper bounds to optimize capital reserves have not been mentioned or explored in the published literature related to private capital cash flows. The closest is Meads et al. (2016) that highlights the importance of effectively managing capital reserves to avoid defaulting on capital calls. The authors consider historical quantiles of cumulative paid-in capital by fund age in the cross-section of funds but do not mention an MPC-like concept for future capital calls. Their analysis focuses on finding an optimal division of the committed but uncalled capital between two different investment vehicles: Treasury securities and the public markets. For a detailed literature review on general cash-flow modeling we refer readers to our previous article (Jeet and O'Shea 2018) in which we explored advanced modeling choices for predicting the *expected* size of future capital calls and showed a large improvement over the approach of Takahashi and Alexander (2002).

4 Methodology

In this section we detail the methodology used to compute MPC. As mentioned in the previous section the methodology is what is often described as a historical methodology in the risk literature (Mina and Xiao 2001). The methodology has two components. First we describe how a PDF of cash flows is generated (sampling, below); this section is relatively novel. Second, we compute various statistics based on this PDF (statistics, below); this section is completely standard.

At a high level our approach is straightforward. Given a large database of historical fund cash flows, and given a fund we wish to analyze, we can find cash flows from funds which, at the time of the observation, were similar to our fund. This generates a pool of observations from which we can randomly sample a certain number of observations per period. Given a portfolio of funds we can carry out this procedure for each fund in the portfolio and then aggregate the cash flows per sample. The result of this procedure is a PDF of portfolio cash flows drawn from the historical behavior of funds similar to those in our portfolio.

Next we describe the methodology more formally.

Sampling In order to define our methodology we first choose an *analysis horizon* (such as a quarter, or a year). A period of time with length equal to the analysis horizon will be referred to as a *period*; we will index periods by t. We also choose a particular such period, the *analysis period*, t_0, to be the period for which we want to make predictions. For each period we also define the *number of samples* we will draw, n_t (this number will be relatively large for recent periods and will decline to zero at some point in the past). Next we choose what we consider to be *cash flows*[4] in this methodology by defining either $X = C$ (cash flows are just contributions) or $X = C - D$ (cash flows are contributions net of distributions).

We need a universe of funds for which we have attributes and cash flows; we will index these funds by i. Each fund has a set of, potentially time-dependent, attributes (such as subclass, size, age, uncalled capital); we denote the attributes of the ith fund at the end of the tth period by $A_t^{(i)}$. Similarly we denote the cash flows of the ith fund during the tth period by $X_t^{(i)}$.

Our methodology needs a notion of similarity between funds. This is defined in terms of fund attributes (for example perhaps we define two funds to be similar if they belong to the same subclass and if their ages differ by less than six months). In general, if a fund at time t (with attributes A_t) is considered similar to another fund at time t' (with attributes $A'_{t'}$) then we write

$$A_t \sim A'_{t'}.$$

Note that this notion of similarity, $\sim$, should be thought of as a (compound) parameter of the methodology, and in fact later we will discuss the trade-offs of different choices.

We start by defining the methodology for a single fund, the *analysis fund*, with attributes A_t. For each period t we sample a number of cash flows from similar funds:

$$S_{t,1}, \ldots, S_{t,n_t} \quad \text{sampled randomly from} \quad \left\{ X_t^{(i)} \mid A_t^{(i)} \sim A_{t_0}, i \text{ is any fund} \right\}.$$

[4] Throughout this document we choose the sign of cash flows to be such that a positive cash flow indicates the flow of funds from the investor into the fund. This is the opposite of the most common sign convention used by investors, but given our focus on capital calls, it is the most natural for this paper.

More generally we will have an *analysis portfolio* consisting of K funds. Corresponding to each of these funds we will generate, as above, samples

$$S_{t,1}^{(k)}, \ldots, S_{t,n_t}^{(k)} \quad \text{for} \quad k = 1, \ldots, K.$$

If these K funds have weight $w_1, \ldots, w_K$ in the portfolio, then we define our portfolio samples as

$$S_{t,j}^{P} = w_1 S_{t,j}^{(1)} + \cdots + w_K S_{t,j}^{(K)} \quad \text{for} \quad j = 1, \ldots, n_t.$$

Finally we pool all these samples $S_{t,j}^{P}$ (for all valid t and all valid j) to form a large set of cash flows which, thought of as a histogram, constitutes the historical prediction for the PDF of the cash flows of our analysis portfolio during the analysis period.

We call the parameters that are needed to define how sampling occurs — namely, the analysis horizon, the analysis period, the number of samples per period, the definition of cash flows, and the similarity measure — the cash flow sampling settings (CFSS).

Statistics At this point we have a PDF of cash flows for our analysis portfolio:

$$\mathcal{S} = \{S_1, \ldots, S_n\}$$

from which we can estimate various statistics. The MPC with confidence α (for example, $\alpha = 95\%$) is the αth quantile of $\mathcal{S}$, namely the number s such that the fraction of elements of $\mathcal{S}$ that are less than s is α:

$$\text{MPC}_\alpha \triangleq \text{smallest } s \text{ such that the fraction for which } \mathcal{S} \geq s \text{ is at least } \alpha.$$

We also compute

$$\text{expected cash flows} \triangleq \operatorname{mean} \mathcal{S},$$

but note that this is the expected cash flows as estimated using the historical methodology. In general this expectation is better estimated using a model-based methodology and in fact is the topic of our previous working paper (Jeet and O'Shea 2018).

Comments on the methodology Like all historical risk methodologies we need a large set of historical observations to ensure that the tails of the PDF are well-sampled (Pritsker 2001). As a result, the parameter n_t above, which dictates how many samples to draw from each historical quarter, and hence determines the relative weight of the recent and distant past, cannot decay too fast. In fact we typically keep n_t constant for a couple of decades and then taper it down to zero quickly. Trying to let n_t decay exponentially would risk reducing the sample size excessively (unless the half-life was made very long).

The requirements on large sample size also affect the similarity measure, $\sim$, defined above. As discussed, we sample from "similar" funds in the past, where we typically define two funds to be similar if they share subclass and have ages that differ by a relatively small amount. In principle one could also demand that funds be similar in terms of uncalled capital, or fund size. The danger, again, is that this could reduce sample sizes excessively. Nevertheless, conditioning on some additional characteristics (especially uncalled capital) seems appealing. One way to do this is to develop a parametric methodology,[5] what we called a model-based methodology above. In such a methodology the distributional assumptions would be explicitly parametrized (for example, contributions could be exponentially distributed with a point mass at zero) and certain parameters would need to be estimated. Such an approach would be less demanding in terms of data and could make conditioning on further fund characteristics feasible. While we think that such a model-based approach is interesting, we will not discuss it further in this paper.

A further question one could ask is why, in the above sampling procedure, we maintain the identity of quarters during portfolio formation, and only then pool samples to form the histogram that we use to compute our risk measures. Indeed, maintaining the identity of quarters is not necessary and dropping this requirement would, in effect, increase our effective sample size (by allowing more mixing). However, there may be temporal

[5]We use the term *parametric* in the sense of parametric statistics. However note that an oft-cited document that describes VaR methodologies (Mina and Xiao 2001) uses the term *parametric* differently from us (they use it to mean that pricing functions are linearized). To avoid confusion we refer to a parametric methodology as a *model-based* methodology.

correlations between cash flows, namely there may be periods when capital is called more rapidly and other periods when it is called more slowly, and furthermore this could be the same for the entire market. By keeping the identity of quarters we are preserving this correlation. Not doing so could potentially lead to underestimating risk (i.e., underestimates of MPC).

5 Model Predictions

In this section we explore the typical behavior of the methodologies described in the previous section as we vary various fund and portfolio attributes.

Our methodology requires choosing values for a small set of parameters, the CFSS. We start by listing the values of those parameters used in this section:

Definition of cash flows Mostly we do not apply netting,[6] so cash flows are defined to be *contributions*. However at times we will contrast predictions of contribution-only cash flows with predictions of net cash flows.

Analysis horizon We make predictions for a period of length one *quarter*.

Similarity measure Our historical methodology requires a similarity measure. In this section we define similar funds to be those that

- belong to the same *subclass*, and
- have ages that differ by less than *three months*.

Samples per period We use 30 samples per period for recent quarters (until 2003) then tapering down to zero (by 1993).

Analysis date We analyze our portfolio as of 2017 Q3.

Next we define our portfolio. It contains funds from five subclasses (buyout, venture capital, real estate, primary equity FoFs, and secondary equity FoFs). For each subclass the portfolio contains eleven funds of ages $0, 1, 2, \ldots, 10$. Thus the portfolio contains a total of 55 funds. We commit one unit of capital to each fund, so our total commitment is 55. In general in the following discussion we will always predict the unscaled cash flows from this portfolio; in particular we will not be reporting cash flows as fraction of total commitment.

Cash flows PDFs Figure 2 contains various histograms that illustrate the result of applying our methodology to our portfolio. We start our analysis by looking at the forecasted behavior of the entire portfolio (figure 2a). The colored vertical lines show the value of various measures computed for the entire portfolio: the prefix ("C." or "N.") indicate whether the statistics are computed using just contributions or contributions net of distributions. The suffixes ".90", ".95", ".99" denote various confidence levels for MPC and ".exp" denotes the expected cash flow. Immediately one can see that MPCs are significantly greater than the expected cash flows; for example C.exp is a bit less than 1.5, while C.95 is about 2.25. The benefit of using net cash flows as opposed to just contribution grows as the portfolio contains more funds (this is because it is a diversification effect; distributions are very volatile, but in the presence of enough funds they become somewhat more predictable and can offset more of the contributions). We can see that the N.95 is significantly lower than C.95. Finally, this plot also shows that the N.exp is negative (i.e., the portfolio is generating more distributions than absorbing contributions) which further illustrates how inappropriate expected cash flows are for setting reserve requirements. In contrast, all the net MPCs (N.90, N.95, and N.99) are positive, as expected.

Next, in figure 2b we break out cash flow PDFs for each subclass. Aside from the cash flows becoming smaller (since each histogram corresponds to commitments to a single subclass), we can also see that the benefit arising from using net cash flows, relative to just contributions, has declined. This can be seen by comparing (say) N.95 (dashed green) with C.95 (solid green); in the top-level PDF the net MPC was about half of the contribution-only MPC, while in these single-subclass portfolios the two are much closer.

[6] By *netting* we mean treating the cash flows that need to be forecast as the contributions net of the distributions in that period. This is in contrast to ignoring the distributions and defining the cash flows to be just the contributions. See the start of section 4 for further details.

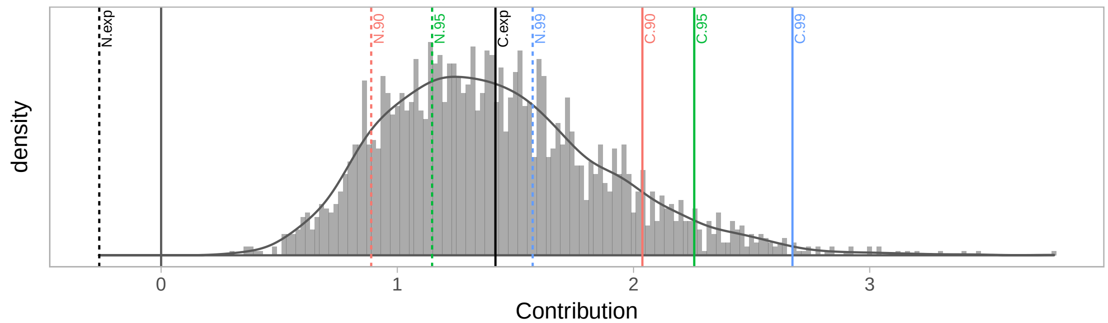

(a) Top level (portfolio) cash flow PDF

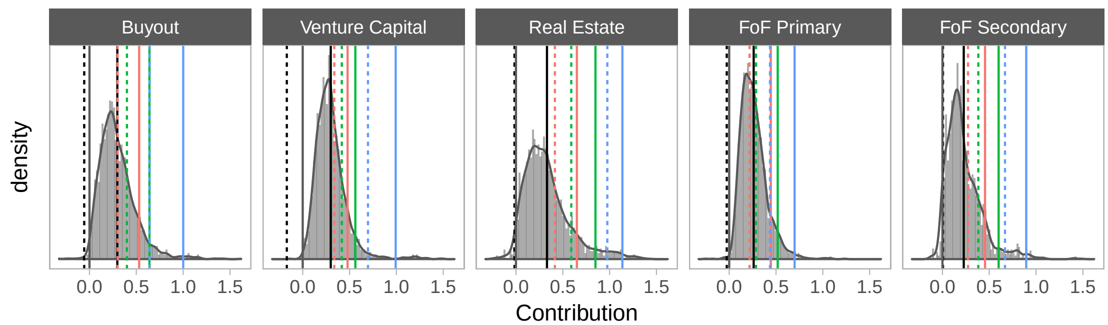

(b) Per-subclass cash flow PDFs

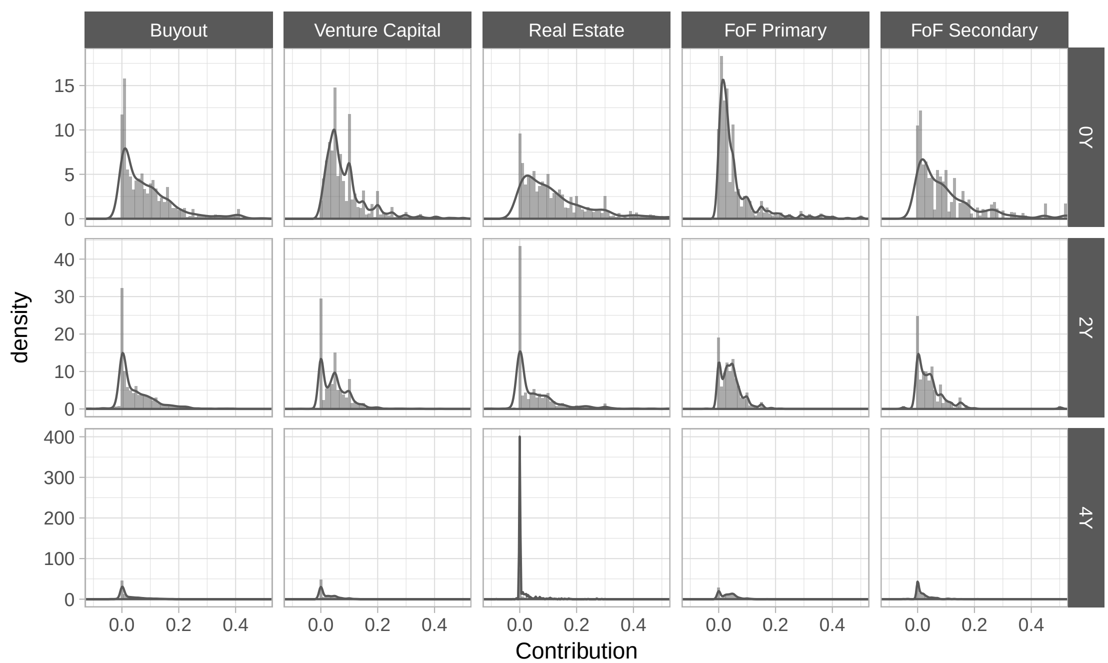

(c) Fund-level cash flow PDFs for a sample of three fund ages (out of a total of 11)

Figure 2: Cash flow PDFs at various levels for a sample portfolio of 55 commitments
The horizontal axis represents unscaled cash flows, namely cash flows arising from a commitment of one unit to each fund

Finally in figure 2c we plot PDFs for a subset of the funds. In this plot it is clear how the younger funds (say the 0Y funds) are calling more capital, but there is a significant uncertainty about just how much they will call (thus in the figure the histograms appear to have an exponential tail to the right). In contrast the older funds call much less and have a larger spike at zero (corresponding to the fact that in many quarters they call precisely zero).

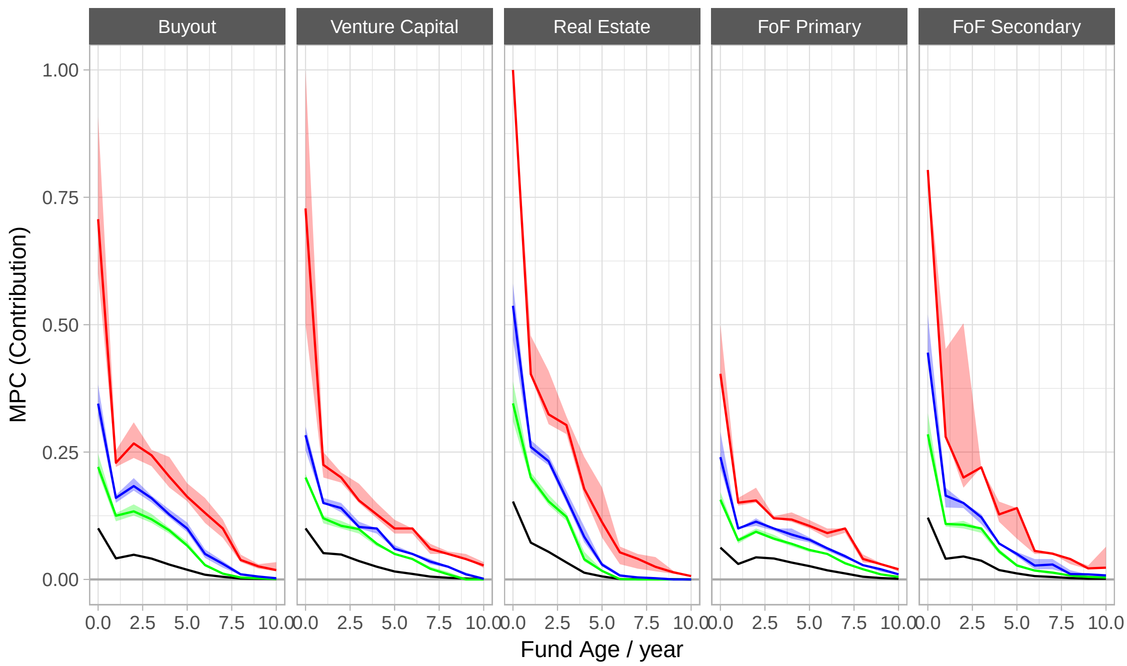

Figure 3: One-quarter MPCs and expected cash flows of a single fund, as a function of fund age
In the above figure the 90%, 95%, and 99% MPCs are plotted in green, blue, and red. The expected cash flows are plotted in black. (The shaded ribbon indicates a 95% confidence interval on the quantile estimators, but should be interpreted as an approximate lower bound on errors.)

MPC as a function of age Figure 3 shows how the MPC of a single fund varies as the fund ages. In that figure we can see various stylized facts about fund behavior. First, note that for all subclasses very young funds have a very sharp MPC peak at inception (zero age); this is partially the nature of how these funds call capital, but is also partially the result of the Burgiss Manager Universe (BMU) defining the inception of a fund to be when its first cash flows occurs. Real estate has very large MPCs early on (especially for high confidence levels) since many such funds call most of the capital very early on. In contrast primary equity FoFs have gentler spikes early on, as expected since the FoF itself must first commit capital before any capital calls can be passed through to the downstream LPs.

Contributions versus net contributions Figure 4 compares contribution-only with net contribution 95% MPCs for a single fund from each subclass. Not surprisingly the net MPC is smaller than contribution-only MPC, however for regular funds (not FoFs) the difference is marginal. For FoFs using net cash flows makes a more significant difference. The reason for this is that distributions are much less predictable (in the sense of being much more disperse). This means that while the expected net cash flows may (for old enough funds) become negative (i.e., become a net distribution) the MPC is unable to leverage those distributions since they are too disperse. Since FoFs are, in effect, a portfolio of funds their distributions are diversified and hence less disperse; consequently they offset the contributions in a more predictable way and hence bring down the MPC to a greater degree.

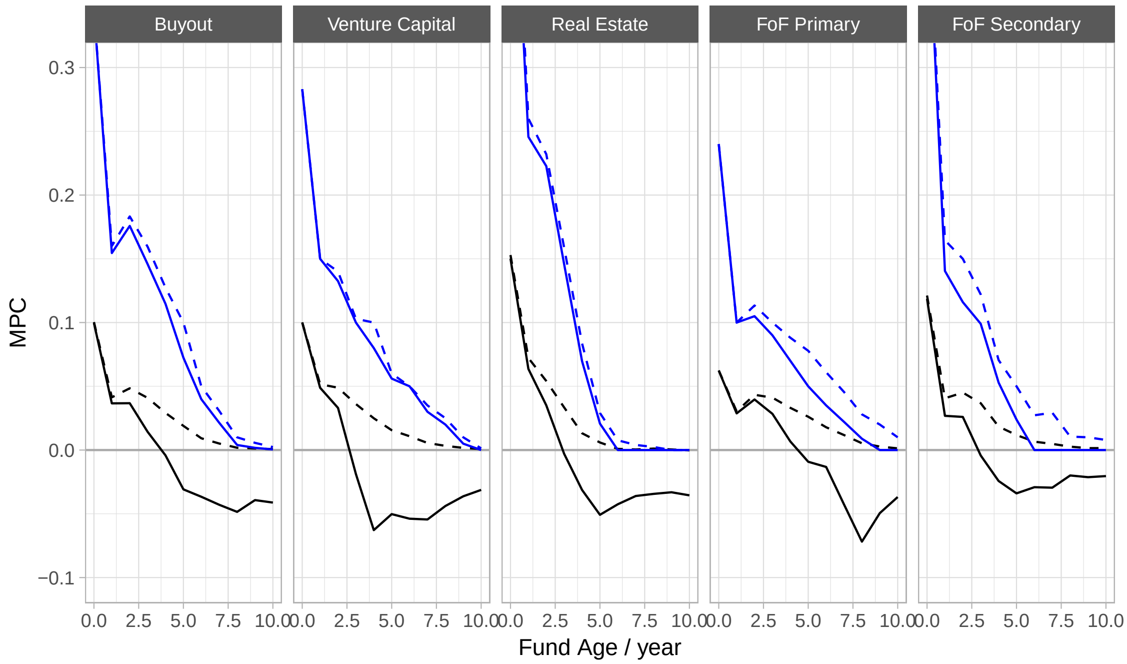

Figure 4: One-quarter 95% MPCs and expected cash flows of a single fund, as a function of fund age
In the above figure the 95% MPCs are plotted in blue. The expected cash flows are plotted in black. The dashed lines are the contribution-only MPCs, while the solid lines are net cash flow MPCs.

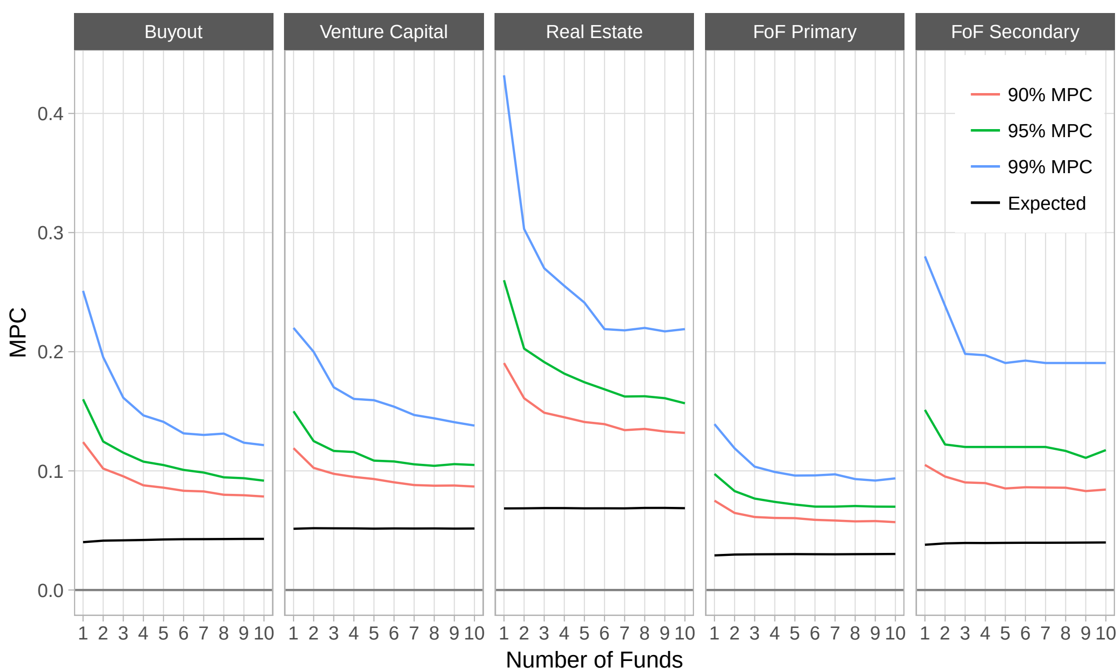

Figure 5: One-quarter cash flow statistics for an increasingly diversified portfolio
The plot shows the expected cash flows and MPCs (at various confidence levels) of a portfolio of distinct funds, each of age 1 year, and drawn from the same subclass (indicated by the facet label). The total commitment is kept at 1 but it is distributed over an increasing number of funds.

MPC as a function of number of funds Figure 5 plots how the MPC of a portfolio of funds varies as the number of funds in the portfolio is increased. In this figure we always have a fixed commitment (of 1.0) however we assume that we have divided that commitment equally among n funds of age 1Y (as of analysis). These funds, although identical (in terms of age and subclass), are distinct and hence diversify each other. Start by noting that the expected contributions (black line) are constant. This is because the expectation of the portfolio is the sum of the expectations of its constituents.[7] Again this serves to illustrate unsuitability of cash flow expectations for the purposes of setting capital reserves. What one would expect is that a single fund would be much less predictable than a portfolio of such funds, and that one would need more reserves. Indeed, plotting MPCs, this is exactly what we see. Furthermore, we see a clear difference between FoFs and funds. The decline in the MPCs of portfolios of FoFs tapers off at about 2–3 FoFs for secondary FoFs and at about 3–4 FoFs for primary FoFs. In contrast portfolios of funds continue to exhibit diversification (declining MPCs) for larger numbers of constituents in the portfolio. This is exactly what one would expect given that FoFs are already a portfolio of funds.

An interesting feature of figure 5 is that the MPC (for any confidence level) does not appear to be converging to the expected cash flows (the black line). It might at first seem that this should occur (as a consequence of the central limit theorem). The observed behavior is explained by two effects. First, the convergence is relatively slow (the gap should be approximately inversely proportional to the square root of the number of funds). Second, there is a common factor driving all contributions. As a result of this factor contributions are partially cross-sectionally correlated. In fact, this is precisely why in our methodology we chose to maintain quarter identity during portfolio formation, and only pool at the very end. We feel this is an important feature of this methodology.

6 Backtesting

It is important to understand how accurate our estimates of MPC are in practice. For this purpose we performed a backtesting analysis for individual funds as well as portfolios of funds. The backtesting procedure we employ is standard.[8] For a given fund and a quarter we predict a PDF for the next quarter's contributions using the methods in section 4. We use the same CFSS listed at the start of the section 5, except we used a similarity measure with an age window of six month (instead of three) and we only allowed data two quarters before the analysis quarter into the PDF prediction. For example if we are predicting MPCs for 2017 Q3 we do not use data from both Q3 and Q2 of 2017. The reason for this is to mirror the typical data-production delays due to lagged reporting in practice.

6.1 Fund-Level Backtests

Figure 6 displays a representative fund-level backtest that ran for 40 quarters and compare the performance of several MPCs with the observed contribution following their prediction. One can see that whenever a contribution was made, its magnitude almost always exceeded the expected contribution (displayed in pink). The median estimate (displayed in red) did even worse because it was almost always zero due to zero-inflation (ZIF) in contributions. Clearly neither the median nor the expected contributions serve as a reasonable estimate of needed capital reserves. For this particular backtest the 90% MPC would have done much better. In principle a 90% MPC will be exceeded, in long run, only 10% of time; similarly a 95% MPC will be exceeded only 5% of time. Note too that for this backtest the highest value of 95% MPC would have been just 30% of the commitment (near inception).

For our fund-level backtesting analysis we ran more than three thousand backtests; one per fund for funds in three major subclasses in the US private capital market: buyout, real estate, and venture capital. These backtests generated a predicted PDF for each observed contribution across all funds and quarters. We summarize this data in two different ways so that it is easy to comprehend, visualize, and gain insights.

One simple way to summarize an observed contribution and its predicted PDF is to compute a binary-valued measure called *excession*. An excession is assigned a one if the observed contribution exceeds an MPC (say 95%), and a zero otherwise. This way the theoretical PDF of excessions would be binomial. For instance,

[7] More formally, the expectation operator is linear.
[8] For more details on our backtesting set up we refer readers to appendix A.

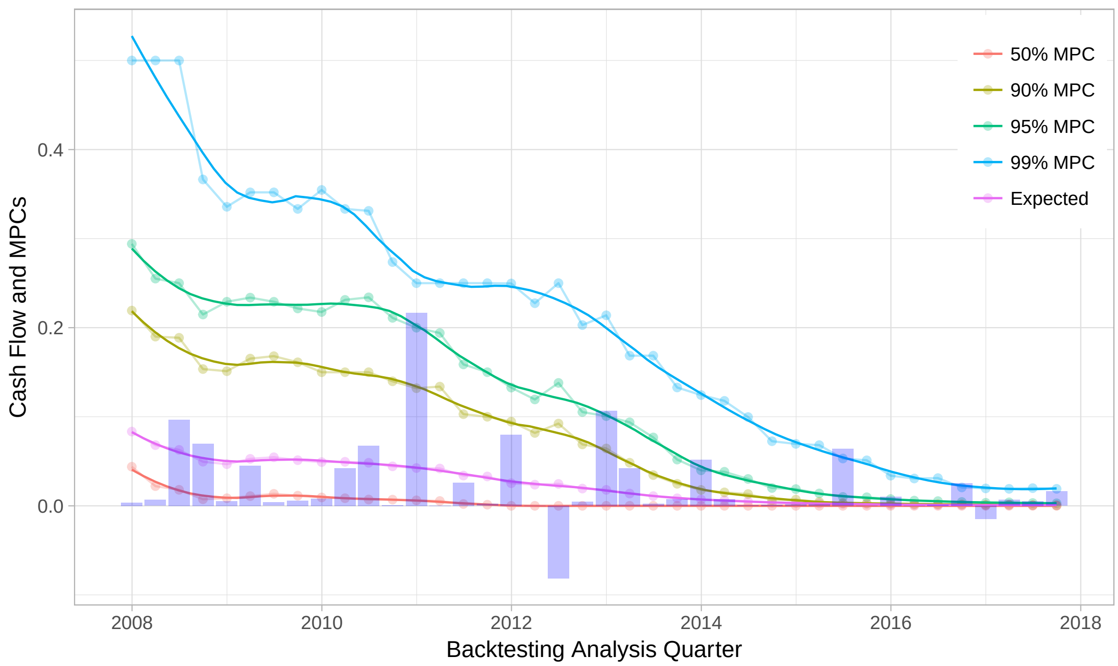

Figure 6: A sample fund-level backtest
The blue bars represent the actual quarterly contributions and solid lines are smoothed time-series MPCs. Note that the negative bar is, in fact, a contribution but represents returned capital.

say in a given quarter there are 100 funds, for 95% MPC we would expect Binomial$(100, 0.05)$ excessions.[9] Similarly in the temporal direction, say contributions of a single fund over 40 quarters, for 95% MPC we would expect Binomial$(40, 0.05)$ excessions.

Another more intelligent way to summarize this data is to compute a quantile rank of the observed contribution from the predicted PDF. For sufficiently large number of observations we can expect these quantiles to be roughly *uniformly* distributed between zero and one.[10]

Figures 7 and 8 plot the PDF of excessions observed both in temporal and cross-sectional dimensions. The theoretical PDF is also included for comparison (in red). One can note that the two PDFs match very well in the cross-section of funds but not so much in the temporal dimension. This is expected and can be explain by two factors. First, the temporal dimension has fewer observations (90 quarters) as opposed to the cross-sectional dimension that has several hundred observations (funds) in each subclass. The second, and a deeper reason, is the business-cycle variability across time that is inherent in the observed data but it is missing from the theoretical model.

Figure 9 display the data of figure 7 in a time-series fashion. This is useful because this plot can detect the effectiveness MPCs over the long-run business cycles and during broad-market events. We observe that most data points stay close to their expected values (shown in black horizontal lines at 10%, 5%, and 1%) and within the two standard deviation interval (gray colored area) but there are some large deviations (venture capital in particular) around times of market crises. There appears to be an absence of excessions during the broad-market crises. It is possible that most of the time variations in the frequency of excession (blue dots falling outside the gray area) are due to continuously changing market conditions before, during, and after the two major broad-market events: dot-com crash (DTC) and global financial crisis (GFC).

Figure 10 displays the histogram of quantile ranks produced in our fund-level backtesting analysis. Given that the quarterly fund-level contribution data is very noisy, the histogram of these quantile ranks is remarkably close to uniformly distributed. The low density for the lower quantiles, and therefore high density for the higher quantiles, is perhaps due to the some interplay of three issues: the noise in the contributions data, the

[9] A note on notation: Binomial(N, p) is an integer-valued random number drawn from a binomial distribution with mean $N \times p$.
[10] Appendix A provides a good discussion on comparing excession and quantile rank.

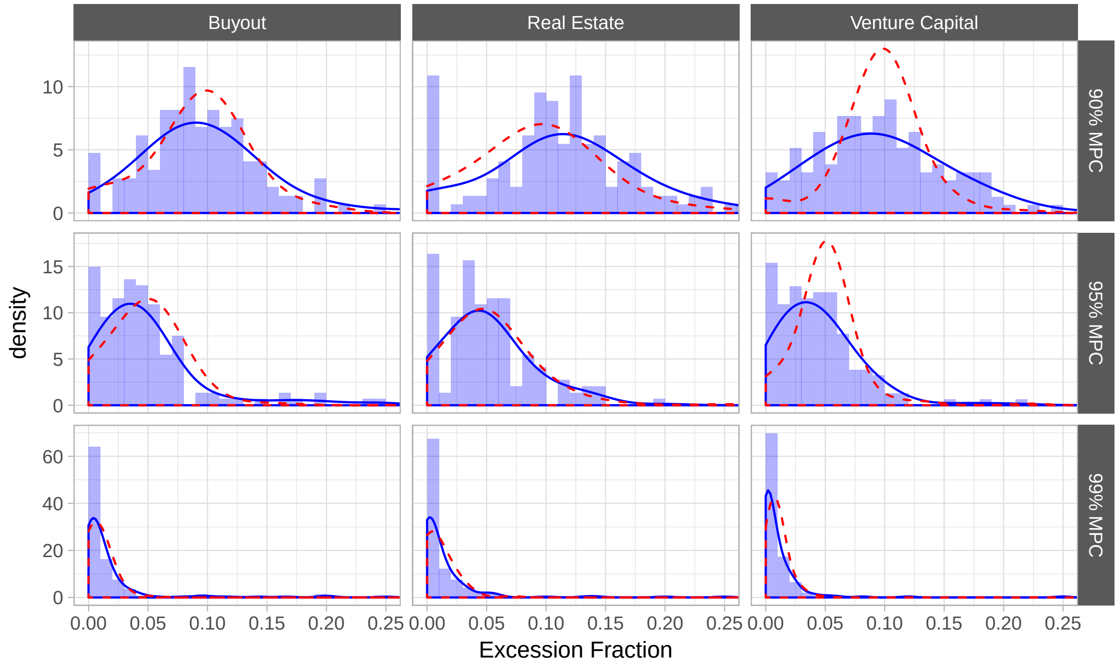

Figure 7: Comparison of theoretical and observed binomial PDF across quarters
The red line represents theoretical and the blue line (and bars) represents the observed PDF.

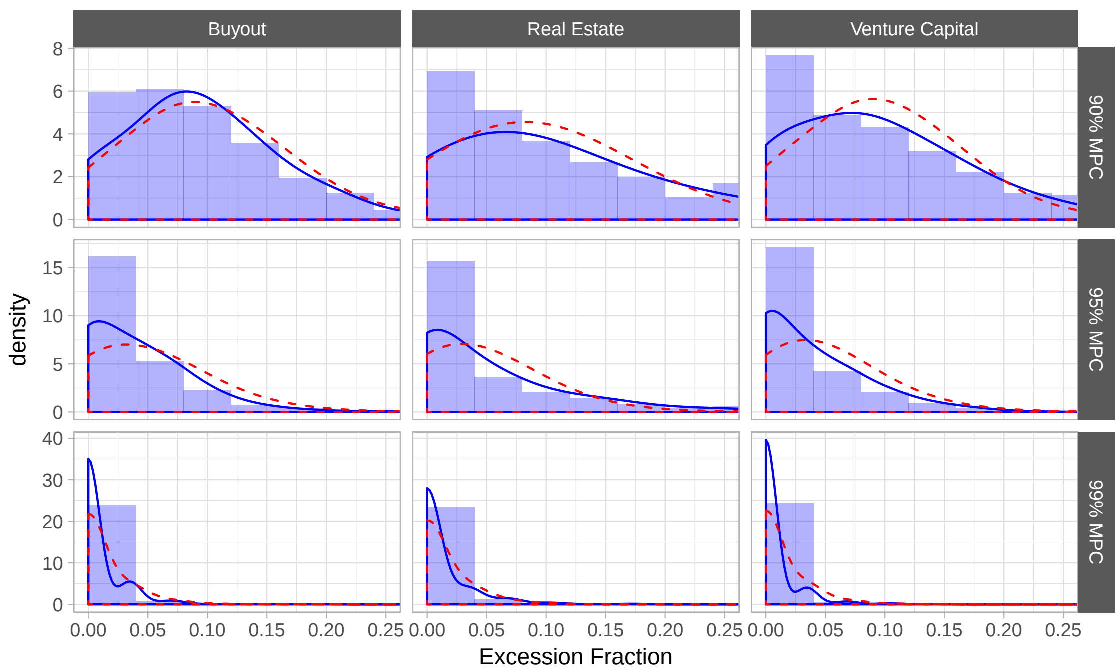

Figure 8: Comparison of theoretical and observed binomial PDF across funds
The red line represents theoretical and the blue line (and bars) represents the observed PDF.

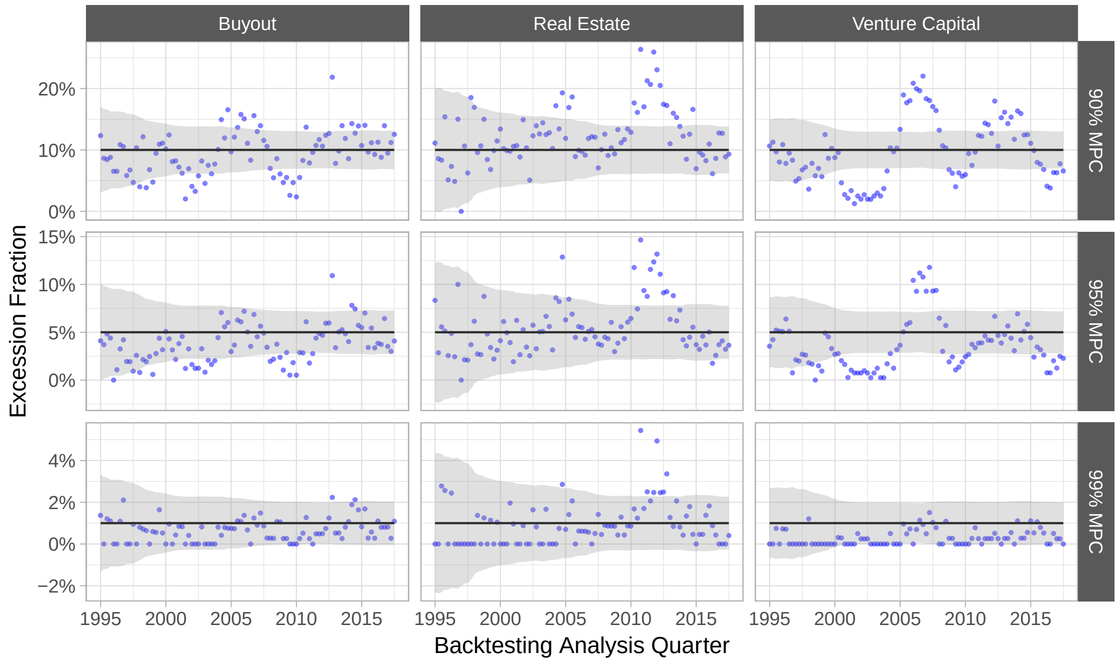

Figure 9: Average excession across time
This plot is useful to detect clustering of (or absence of) excession events across time.

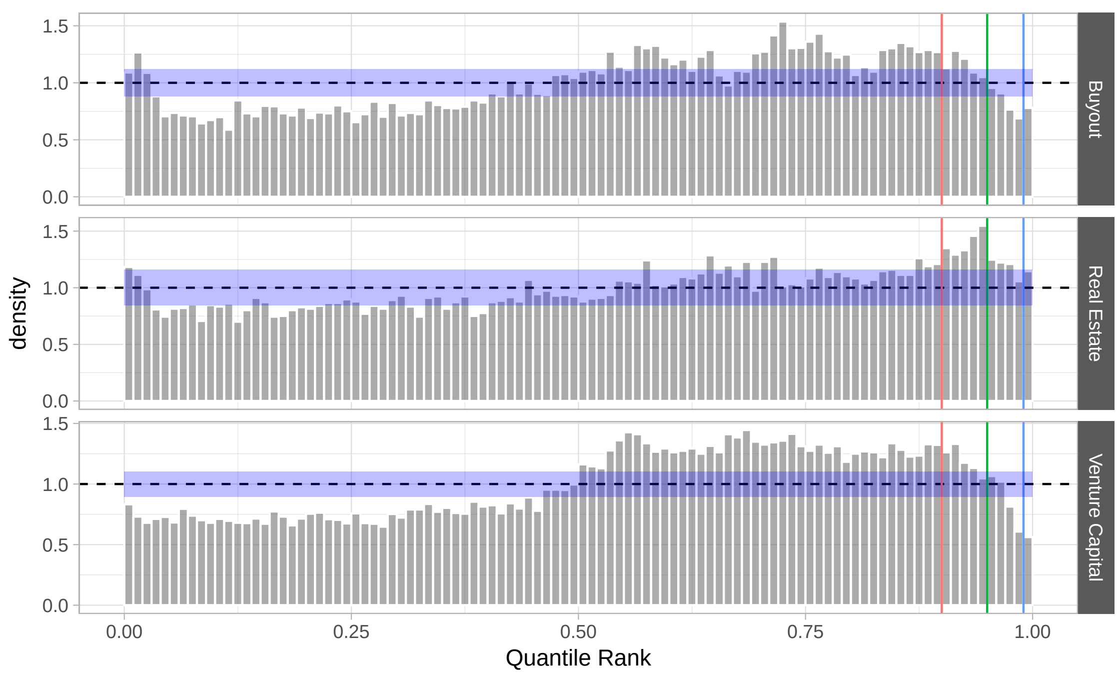

Figure 10: Histogram of quantile ranks of quarterly fund-level quarterly contributions
The theoretical expectation of this histogram is the uniform PDF represented by a dashed line. The shaded area in blue covers the 95% confidence interval for the height of each quantile bar. The three vertical lines in red, green, and blue represent 90%, 95%, and 99% quantiles respectively.

ZIF problem, and the limitations of the historical-simulation methodology. The expectation that quantile ranks be distributed uniformly does not take into account ZIF. The presence of ZIF distorts the PDF of quantile ranks and requires ZIF-sensitive quantile-ranking methods, see appendix A.1 for more details.

6.2 Portfolio Backtests

Next we move to backtesting of portfolios of private capital funds. For this analysis we build random portfolios with a simple policy of investing in one fund per vintage from 1995 to 2017. In each vintage the selected fund could come from either buyout, venture capital, or real estate subclass. All selected funds are weighted equally to form a portfolio. Our analysis is based on 1000 such portfolios sampled randomly. Figure 11 displays backtest of one such portfolio. Note that a portfolio of funds almost always makes contributions every quarter. Similar to what we saw in fund-level backtest that the *expected* contribution limit is frequently exceeded by the actual amount of contributions following prediction. For the particular portfolio shown in figure 11, the 95% MPC was exceeded only three times over a period of 22 years. What is more interesting is that portfolios enjoy significant diversification benefit as well, which means that capital reserves are shared among the funds that are actively calling capital. A portfolio that invests in one fund per vintage may have, in the long run, 5 to 7 active funds[11] that can potentially call capital, yet the capital reserves requirement set by 95% MPC is less than the commitment into a single fund.

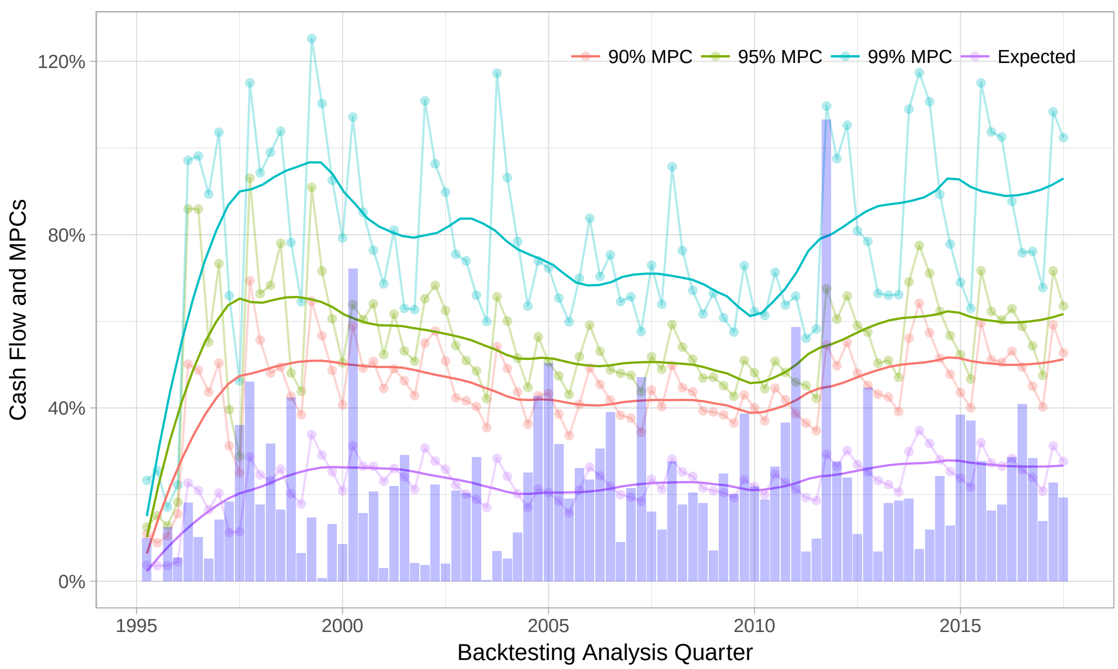

Figure 11: A sample portfolio-level backtest
The blue bars represent the actual quarterly contributions and solid lines are smoothed time-series MPCs.

The observations we made above are specific to the portfolio selected for figure 11. In order to test their statistical significance we plot summary information from these backtests in figures 12 and 13. Figure 12 compares expected contributions with three MPCs levels: 90%, 95%, and 99%. We plot the density of excession magnitude and color the portion (in red) where the excession is positive, i.e., the percent of times an observed contribution exceeded the stated MPC level. One can see that across all contribution events, observed for 1000 portfolios over 90 quarters, *expected* contribution was exceeded at the rate of almost 50% but the MPCs correctly exceeded only at the rate that is pre-defined, i.e., 5% for 95% MPC and so on.

Figure 13 plots the histogram of the quantile ranks produced in our portfolios backtesting analysis and as is expected the histogram roughly matches the uniform PDF. Unlike the fund-level contributions data (shown in

[11] Since most funds call most of the committed capital within first 5 years.

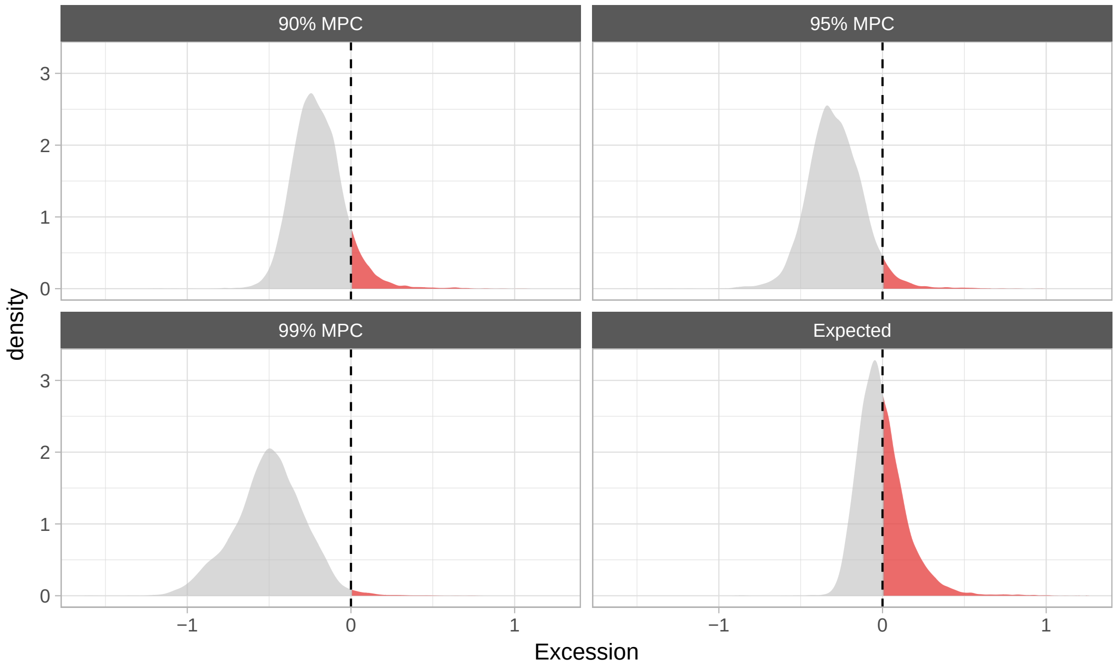

Figure 12: Density plots of excession magnitude
The area on the right-hand side of the dashed-line roughly represent the probability of exceeding the underlying MPC.

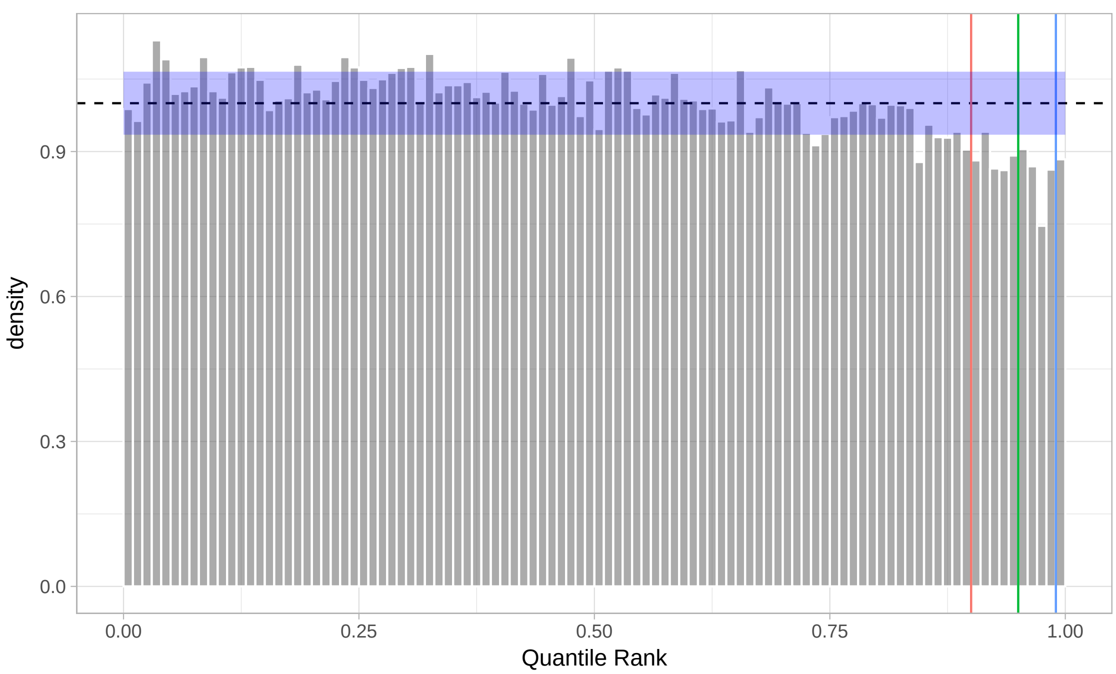

Figure 13: Histogram of quantile ranks of quarterly portfolio-level quarterly contributions
The shaded area in blue covers the 95% confidence interval for the height of each quantile bar. The three vertical lines in red, green, and blue represent 90%, 95%, and 99% quantiles respectively.

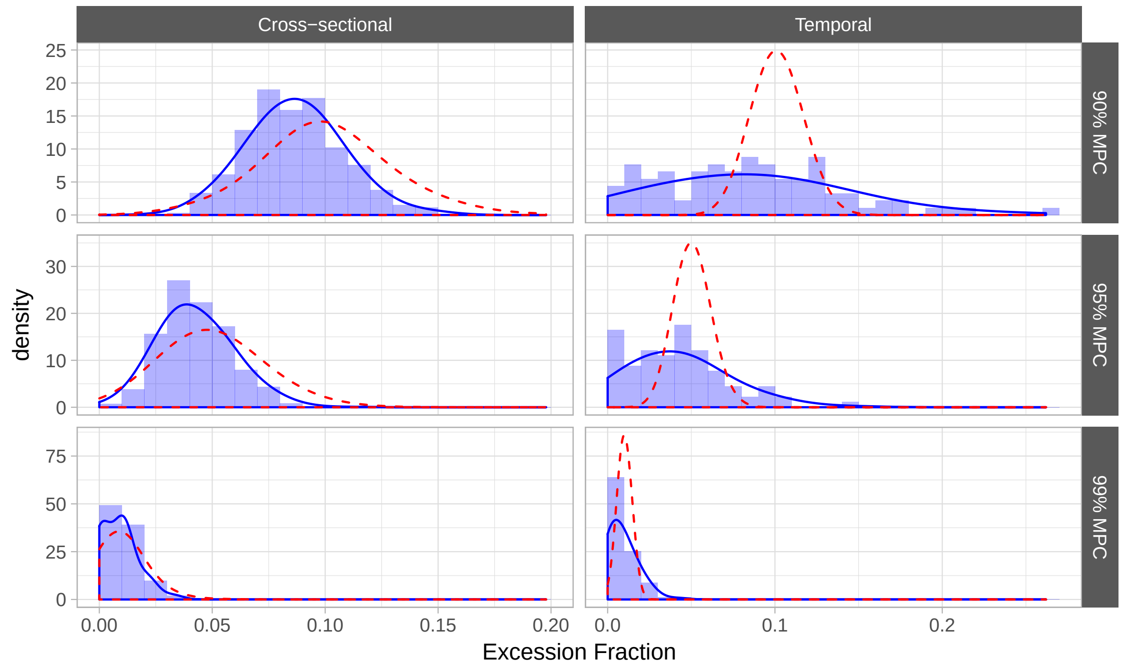

Figure 14: Comparison of theoretical and observed binomial PDF of excession events
The red line represents theoretical and the blue line (and bars) represents the observed PDF.

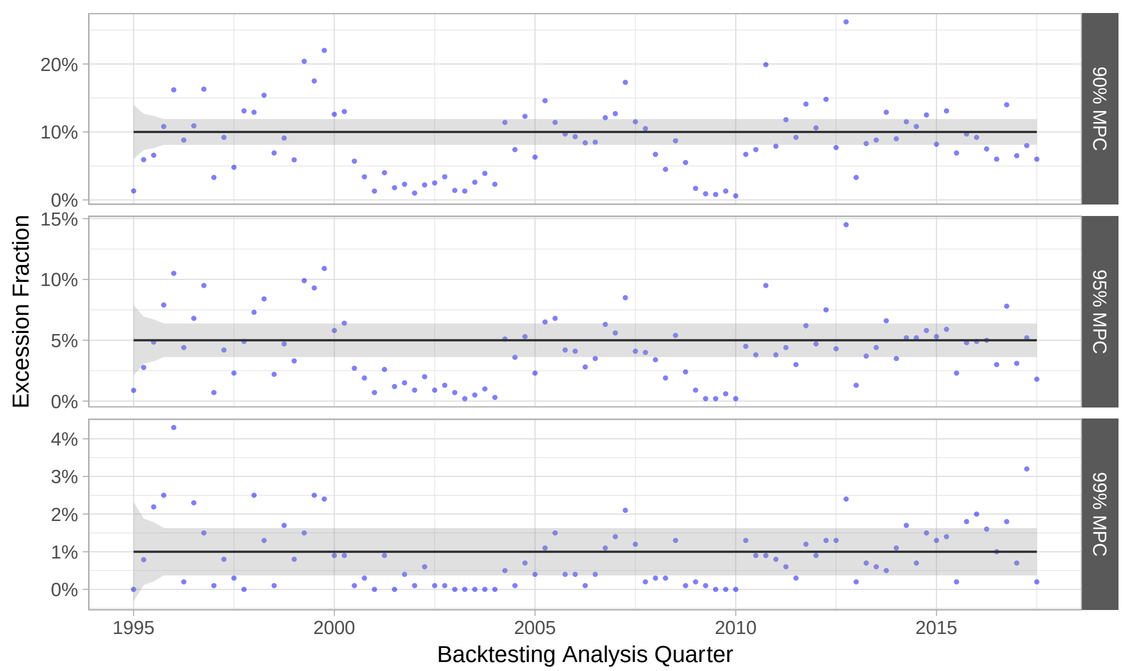

Figure 15: Average excession across time
This plot is useful to detect clustering of (or absence of) excession events across time.

figure 10) the portfolio-level contributions data are relatively less noisy and do not have the ZIF problem. This explain why the histogram in Figure 13 is much closer to the theoretical uniform distribution (represented by dashed black line) than those in figure 10. Although looking at the higher MPCs (95% onward) in Figure 13, it seems they were overestimated. We think this is perhaps due to the limitations of historical-simulations methodology and reporting delays.[12]

Figure 14 plots the PDF of excession observed both in temporal and cross-sectional dimensions for portfolio cash flows, the theoretical PDF is also included for comparison (dashed line in red), similar to what we saw in the fund-level backtesting analysis. The empirical and theoretical PDFs are generally in agreement but with some mismatch in the temporal dimension perhaps, again, because it has fewer observations and the theoretical model does not capture the business-cycle variability inherent in the observed data. Finally figure 15 plots the mean excession in a time-series fashion to detect any clustering (or absence) of excession events during specific period. Once again, there appears to be an absence of excessions during the broad-market crises, specially the GFC.

7 Conclusion

This paper is concerned with making *risk predictions* about future contributions. A risk prediction is one where what is predicted is not the value of some random variable (such as the total capital calls arising from a portfolio) but the distributional properties (i.e., the PDF) of that variable. Once such a PDF is predicted it is straightforward to calculate various measures, such as quantiles, that serve as probabilistic upper bounds on the contributions in the next period. For example, the 95% MPC is an amount such that the contributions in the next period will be less than that amount with probability 95%. These quantiles are *actionable* numbers for investors. For example, investors that maintain reserves at the 95% quarterly MPC can expect to only need funds in excess of those reserves once every five years; maintaining reserves at 99% MPC should result in this happening once every 25 years. In contrast, actual contributions will exceed expected contributions roughly twice a year.

We propose a methodology based on historical sampling for estimating the PDF of these contributions, which in turn allows us to estimate MPCs. This methodology takes into account, in a natural and quantitative way, many expected effects including: the age-dependent rate at which funds call capital, the variation across subclass in how funds call capital, the diversification benefit arising from portfolios that have multiple funds, and the reduction in MPC if contributions are offset against distributions in each period. After illustrating these effects, we also backtested the performance of the risk model and found it to perform remarkably well. The backtesting was carried out against both funds and random portfolios and by simultaneously looking at multiple confidence levels.

In summary, historically-estimated MPCs seem to be accurate, quantitative, and actionable measures of much contributions an investor can expect in the next period. Furthermore MPCs capture many complex interactions between the cash flows of private capital funds, including their age dependence and the degree that they diversify among each other.

[12] In general backtesting only uses data strictly before the observation it is trying to predict. However we skip an additional quarter of data since we think that on account of reporting delays this better reflects the historical data that would be available to any model making predictions.

References

Jeet, V. and L. O'Shea (Jan. 2018). *Modeling Cash Flows for Private Capital Funds*. Working Paper 4. The Burgiss Group LLC.

Meads, C., N. Morandi, and A. Carnelli (2016). "Cash Management Strategies for Private Equity Investors". In: *Alternative Investment Analyst Review* 4 (4), pp. 31–43.

Mina, J. and J. Xiao (2001). *Return to RiskMetrics: The Evolution of a Standard*. Tech. rep. RiskMetrics Group.

Pritsker, M. G. (2001). *The Hidden Dangers of Historical Simulation*. Working Paper. The Federal Reserve Board.

Takahashi, D. and S. Alexander (2002). "Illiquid Alternative Asset Fund Modeling". In: *The Journal of Portfolio Management* 28.2, pp. 90–100.

Zumbach, G. (2007). *Back testing risk methodologies from one day to one year*. Tech. rep. RiskMetrics Group.

A Backtesting Risk Methodologies

In this document a *risk prediction* is a prediction of the distribution (PDF) of a random variable (rather than just its value). A risk methodology is a technique for generating these forecasted PDFs. In general these fall into two classes: model-based (where the PDF belongs to some family and the model's job is merely to estimate the parameters for an instance from that family) and historical (where the PDF is derived non-parametrically from historical data).

A common approach to backtesting a risk methodology is to derive a tail measure, such as the 95th percentile, and then count *excessions*, namely observations that exceed the tail measure. A good model, in this case, will have an excession rate of about 5%. Note the excession rates that are too high or too low indicate deficiencies in the risk model. A disadvantage of this approach is that it focuses on a single confidence level. If one were interested in a different level (such as 99% MPC) then the entire backtest would need to be repeated. In addition this approach discards information from "near misses" since one can only count excessions, there is no such thing as a "small" excession or an "almost" excession.

These problems can be avoided by instead focusing on a different variable, namely transforming each observation into its quantile rank via the predicted PDF (Zumbach 2007). For example, suppose that at each period t we observe a real-valued random variable X_t. Corresponding to each observation we make a prediction for the PDF of X at t, $\phi_t(\cdot)$. Let $\Phi_t(\cdot)$ be the corresponding cumulative density function (CDF), then $q_t \Phi_t(X_t)$ is a number between 0 and 1. Furthermore, if each X_t is, in fact, distributed according to ϕ_t (so the predictions were perfect) then q_t is a uniform random variable between 0 and 1. Furthermore if ϕ_t is the true distribution of X_t *conditional* on all information up to t then q_t should be temporally uncorrelated.

For example, suppose X_t is a standard normal random variable, and suppose we consider three models (or predictions) for its PDF: normal distributions with mean zero but with standard deviations of 1 (correct), 0.8 (under-predicting risk), and 1.2 (over-predicting risk). The distributions of q_t for each of these models are illustrated in figure 16. The top facet (labeled "correct") shows the result of using the correct model. As

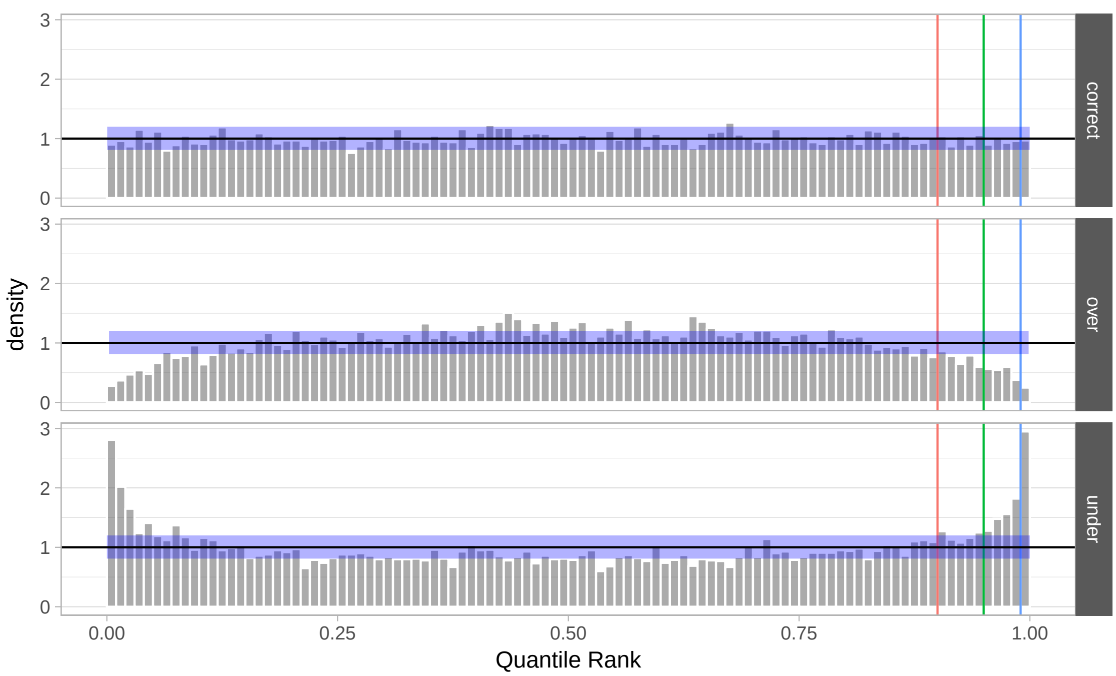

Figure 16: Probability of getting various quantile ranks for several illustrative models
In this figure the data and models are normally distributed with mean zero. The data has a standard deviation of 1 and the models have standard deviations of 1, 1.2, and 0.8. The number of data points is 10^4*. The horizontal ribbon is a 95% confidence interval around the expected density of quantile ranks, assuming the model is correct.*

expected the quantile ranks are approximately uniformly distributed between 0 and 1. Note too that there is a certain amount of natural variation (or noise) arising from the fact that the histogram is based on a finite sample of data. The horizontal ribbon in the figure is a 95% confidence interval on the number of quantile ranks in each bar of the histogram. As can be seen, the correct model stays within the confidence interval about 95% of the time. In contrast the model that over-predicts risk has a smaller number of quantile ranks than expected at both end of the graph (and strays far outside the confidence interval). Similarly, the model that under-predicts risk has too many quantile ranks at each end of the graph and again strays far outside the confidence interval.

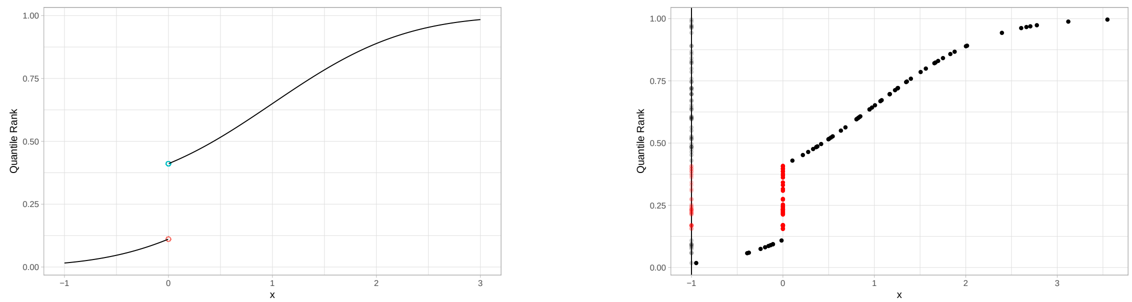

(a) Analytic CDF **(b)** Quantile ranks of samples from distribution

Figure 17: CDF and quantile ranks arising from a distribution with a point mass
The distribution is a normal distribution with mean 1, but with an additional point mass at zero.

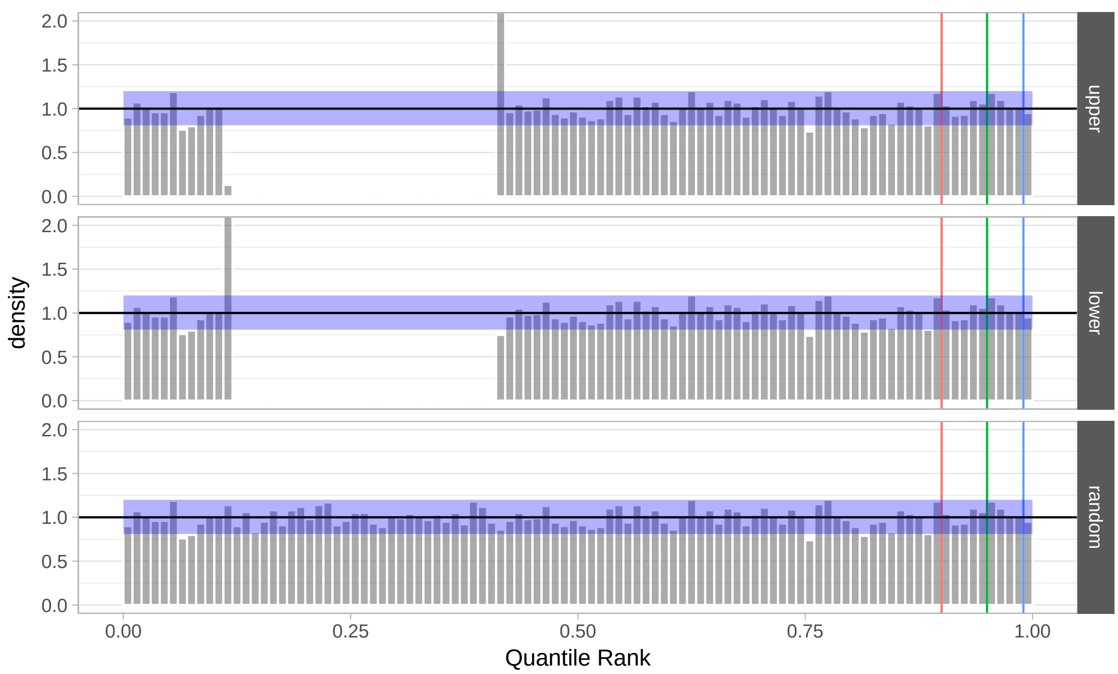

Figure 18: Probability of getting various quantile ranks for a distribution with a point mass
In this figure the real and model distributions are a normal distribution with mean 1 mixed with a point mass at zero. In each facet we map the observations of zero (coinciding with the point mass) to either the maximum quantile (upper) the minimum quantile (lower) or a uniformly at random between the two extremes (random).

A.1 Distributions with Point Masses

The discussion in the previous section regarding quantile ranks assumed that the model distribution is continuous. If the model distribution has point masses[13,14] then the above procedure runs into difficulties. Consider the CDF depicted in figure 17a; it corresponds to a normal distribution with mean 1 and with a point mass at zero. What should the quantile rank corresponding to $x = 0$ be? Choosing either the lower (red) or upper (blue) quantile rank will result in distributions of quantile ranks that are far from uniform. For example, in figure 18 we have plotted the distribution of quantile ranks resulting from these two choices (in the facets labeled 'upper' and 'lower') which clearly show a non-uniform distribution, even though our model is correct! In our backtesting we employ a simple fix for this issue. Instead of choosing the upper or lower quantile we choose a quantile uniformly at random between these two extremes. This choice is displayed in figure 17b where the red points have been distributed randomly between the possible extremes. The facet labeled 'random' in figure 18 shows the result of this fix, which as can be seen makes the quantile ranks uniformly distributed as expected when the model agrees with the data.

B Data used for Model Estimation and Backtesting

All computational results in this paper are based on private capital data in the BMU as of 2017 Q3.[15] Our data consisted of US funds denominated in USD (excluding funds of funds) from three subclasses of funds, namely buyout, real estate, and venture capital. For FoFs analysis we focus on US FoFs denominated in USD that are further categorized as primary or secondary in their market focus. Both funds and FoFs are distributed over a broad range of vintage years from 1980 to 2017. We aggregated dated Contributions and Distributions to the end-of-quarter date for each fund or FoF.

[13]For a purely continuous distribution the probability of observing any particular value (as opposed to an interval) is zero. However if the probability of observing some x is non-zero, then the distribution is said to have a *point mass* at x.

[14]Note that the PDF of private capital cash flows does have a point mass (at zero) since there will be many quarters with no cash flows.

[15]The Burgiss Manager Universe is a research-quality dataset comprised of nearly 40 years of daily cash flows and valuations for over 7,800 private capital funds, representing more than $5 trillion of capital committed across the globe in various Private Equity, Private Debt and Real Asset strategies. The dataset covers a full spectrum of strategies across all geographies, and BMU data is representative of global institutional investor experience because it is sourced entirely from limited partners, avoiding the natural biases associated with other data sourcing models.

111 River Street, 10th Floor
Hoboken, NJ 07030

p: +1 (201) 427.9600
f: +1 (201) 795.9237
w: burgiss.com