YALE CASE SEPTEMBER 23, 2025

Yale SCHOOL OF MANAGEMENT

# Does Recurring Revenue Really Drive Financial Outcomes in Search Fund–Acquired Businesses?

## Recurring revenue is a mantra in ETA land, but our data indicates it might not matter as much as we think

***Daniel Lazier***[1]
***Jacob Thomas***[2]
***A. J. Wasserstein***[3]

Every aspiring search fund entrepreneur aims to acquire a business with recurring revenue. That is gospel in the entrepreneurship through acquisition (ETA) community. Academics, investors, podcast hosts, conference panelists, and full-cycle ETA CEOs all beat the drum of the importance of recurring revenue. We have never seen an ETA private placement memorandum that does not list recurring revenue as an essential target economic characteristic. Typically, recurring revenue is the number one sought-after factor in potential acquisition candidates. Additionally, there are some investors in the ETA jungle who position themselves as recurring-revenue-only investors; that is a fundamental screen for their capital allocation activities. We explored general revenue parameters in our note On the Nature of Revenue, where we paint a picture of revenue quality on a ladder with recurring revenue occupying the top rung.

In spite of all this, revenue properties are just one dimension in the grab bag of vectors that influence financial outcomes for a company. Growth rates, profitability, operational complexity, capital intensity, customer acquisition costs, and leveragability also play roles. In other words, revenue dynamics are an input; they are not an end goal. Entrepreneurs and investors measure success by multiples of invested capital (MOIC) and internal rates of return (IRR). If revenue serves as a bridge to superlative financial results, that is wonderful, but the real aim is not recurring revenue; the objective is eye-popping MOICs and IRRs.

Since recurring revenue is such a touchstone in the ETA ecosystem, we want to isolate that single factor and attempt to ascertain whether or not recurring revenue truly influences financial results. The first step in this analysis is to determine if there is any relationship between a company's percentage of recurring revenue at acquisition and its financial outcomes. While establishing causation is another matter, our main goal is to explore whether any pattern exists between the presence and level of recurring revenue in a firm at acquisition and its MOIC and IRR.

In our analysis, we focus exclusively on MOIC and IRR. Exit multiples, customer retention rates, and margins are not our concern. MOIC and IRR are the metrics that tell the whole story. Everything else may be a relevant stepping stone, but it simply does not matter in our current context. Imagine an ETA CEO boasting about retention rates and exit multiples while delivering a 2.25x MOIC. Who cares? MOIC and IRR rule the day and are the CEO's concise scorecard.

We approach this investigation with curiosity but dispassionately. We do not care what the answers are; we let the data tell the tale. A warning to our readers: Our note will be wonky and peppered with lots of math and statistics. For readers who are not statistics jocks, we have a brief

glossary in **Exhibit 3**. There will be no inspirational stories, no spotlight on winners or losers with drama, personalities, or an impressive saga. This is an empirical, unbiased, data-only assessment.

Although this note is quite dense and technical, we believe it is worth investing the time to read and understand because it challenges an ETA conventional wisdom. Like many things in life, the idea that recurring revenue is remarkably beneficial has been repeated so often and so loudly that ETAers accept it as a truism. Spoiler alert: *Contrary to the widely held belief that there is a strong positive relationship between acquisition outcomes—such as MOIC and IRR—and recurring revenue, our data and analysis show only a weak relationship, which is also statistically insignificant because it is associated with considerable uncertainty*. For the details, read on.

In this note, we will discuss three principal topics (**Figure 1**). We hope that reading it will help students and entrepreneurs better understand the significance, or lack thereof, of recurring revenue in financial results. Let's dig in!

**Figure 1: The roadmap for this note**


## Our methodology

To conduct our analysis, we relied on a valuable data set provided by a single generous investor with decades of search fund experience and results. Specifically, the data set consists of 59 observations (**Exhibit 1**), all of which are traditional search fund projects that have completed the full cycle of raising a fund, acquiring a company, and ultimately exiting. As a reference point, according to the 2024 Stanford Graduate School of Business Search Fund Study, 162 traditional search funds have exited since the model's inception.[4] Researchers always want more data, but our sample of 59 observations is relatively large compared to the 162 firms in Stanford's most comprehensive study. None of the observations in our data set include mark-to-market valuations; they have all been valued by third-party acquirers. Additionally, the helpful investor characterized the presence of recurring revenue in each company by examining the CEO's confidential information memorandum at the time of the initial search fund acquisition. Accordingly, each of the 59 observations has a recurring revenue percentage ranging from zero to 100%. Clearly, there is some subjectivity in the diagnosis, but the recurring revenue percentages at least provide some directional indication of its presence or absence.

We recognize that not all companies with recurring revenue are identical or interchangeable. For example, a software-as-a-service business might have long-term contracts with considerable customer switching costs, both financially and behaviorally. It might also experience positive logo attrition. This is somewhat different from facilities management services companies that have customers on two- to three-year contracts but have 20% annual customer churn due to competition and lower switching costs. The investor

who provided the data did not differentiate between these two types of recurring revenue, despite the natural gradations in quality. A key lesson: No data set is perfect.

We created four bins to classify the 59 companies based on their percentage of recurring revenue at the time of acquisition. Companies with zero recurring revenue form one group (None) with 29 data points, while the remaining 30 firms are divided into three roughly equal bins. Firms with more than zero up to 49% recurring revenue are in our Low recurring revenue bin with eleven observations. Businesses with 50% to 79% recurring revenue are in our Moderate bin with nine observations. Finally, our High category includes enterprises with 80% to 100% recurring revenue, totaling ten observations. So, we have four bins: None, Low, Moderate, and High recurring revenue (**Figure 2**). We will capitalize the bin labels throughout the note so it is clear we are referring to revenue bins.

**Figure 2: The four bins based on our study's 59 observations**

| | None | Low | Moderate | High |
|---|---|---|---|---|
| Recurring revenue range | 0% | >0% to 49% | 50% to 79% | 80% to 100% |
| Bin count | 29 | 11 | 9 | 10 |

For each observation, we also have the ending IRR and MOIC, as well as the EBITDA* multiple paid during the acquisition. Consequently, we conducted analyses to identify patterns between recurring revenue, MOIC, and IRR. Furthermore, we examined the relationship between recurring revenue and entry EBITDA multiple. Our data source requested that we mask the raw numbers in the data set to maintain some semblance of confidentiality. We understood this and present the financial results in two ways.

First, we grouped raw MOIC and IRR outcomes into buckets based on cutoff points that reflect varying levels of success from the perspectives of entrepreneurs and investors. We replicated the MOIC buckets used by Stanford in their biennial search fund report, but instead of grouping MOIC outcomes from 1.0x to 1.9x, we grouped them from 1.0x to 2.9x. Our reasoning is that combining 2.1x outcomes with 4.9x results creates a mismatch, as a 4.9x MOIC is excellent, while a 2.1x MOIC is less impressive. We believe these buckets better reflect the distinction between desirable and undesirable outcomes. The MOIC buckets are as follows: 0 (total loss); 0.1x–0.99x (partial loss); 1.0x–2.99x; 3.0x–4.99x; 5.0x–9.99x; and 10.0x or higher. The IRR buckets are as follows: total loss; partial loss; 0–9.99%; 10.0–19.99%; 20.0–34.99%; and 35.0% or more, where the total and partial loss buckets are the same as the corresponding MOIC buckets. In the figures below, we depict MOIC and IRR buckets with only one digit after the decimal place, but the methodology includes all observations and reflects the categorization in this paragraph. We included separate histograms illustrating how firms are spread across these outcome buckets for each recurring revenue bin.

Second, we created a synthetic index to normalize MOICs and IRRs on a scale from 0 to 100, representing the minimum and maximum values in our sample. The remaining observations were then assigned values based on their relative position between these two extremes. For instance, an MOIC value that is exactly halfway between the minimum and maximum received an index value of 50. While these indices may be less intuitive for readers than raw data, they provide a clearer sense of scale and variability within the data

* Earnings before interest, taxes, depreciation, and amortization.

set while maintaining the original data distribution. Additionally, we transformed the raw MOIC and IRR data into ranked data, which not only obfuscates the actual figures but also reduces the influence of outliers. We will briefly discuss the ranked analyses below.

Creating the indices allows us to generate box plots that display key statistics for the distributions—such as mean, median, and interquartile range—of MOIC and IRR across different recurring revenue bins. The indices also enable us to perform various statistical tests to analyze differences in outcomes among the bins. Additionally, we can estimate correlations and regressions between outcomes and recurring revenue, reporting t-statistics and p-values to determine if these relationships are positive and statistically significant.

Our study is certainly not perfect, but overall, the analyses we performed provide some directional insights into recurring revenue potential in a search fund project. All our data comes from a single investor, which could introduce sampling bias. However, utilizing data from one investor helps us avoid double-counting certain outcomes from a pool of investors. For example, if investor A has investments in companies 1 and 2, and investor B has investments in companies 2 and 3, the results for company 2 would be weighted more heavily (counted twice) in the analysis if we combined anonymized data from both sources. That is not the case for our data set.

Our study is also potentially subject to omitted variable bias because we are singularly examining the relationship between recurring revenue and economic outcomes. We are not performing multivariate analyses or considering other variables that might influence or correlate with financial results. These variables could include industry, acquisition vintage, or CEO demographics. In general, our results could be biased if these variables are related to both outcomes (MOIC and IRR) and recurring revenue. Again, because the data comes from a single investor, that investor's unique selection criteria are likely magnifying the effects of omitted variable bias. As a result, some variables in our analysis may have inflated coefficients, while important effects of other variables may be washed out or dampened by the investor's selection criteria.

We are isolating recurring revenue for a few reasons: It is a key point in the ETA space; we lack complete data for other variables; and we want to keep the analysis simple for ourselves and our audience. We will now shift gears and discuss the results of our work.

## The results

We approached our analysis from multiple angles to gain a comprehensive understanding of how recurring revenue impacts ETA firms. We started with percentage frequency histograms of MOIC buckets for each of the four recurring revenue bins (**Figures 5** and **8)**. We then repeated the histograms for partitions of indexed MOIC. Recall that the raw MOIC buckets correspond to six buckets based on the actual MOIC earned by firms in the sample, while the indexed MOIC normalizes raw MOIC outcomes to range between 0 and 100. Next, we drew box plots for indexed MOIC across the four recurring revenue bins. We then repeated the analysis for IRR. Finally, we performed various statistical tests, examining both indexed MOIC and indexed IRR.

Although it involves a lot of math and numbers, all evidence points to the same conclusion: *We cannot reject the null hypothesis that there is no link between high recurring revenue firms at acquisition and economic results*. Of course, this does not mean that the alternative hypothesis, which suggests a positive relationship between

recurring revenue percentages and MOIC and IRR, is false. Just because we cannot reject something does not necessarily mean that it is true and that the alternative is false. *Overall, our results imply that there is no strong relation between recurring revenue and MOIC and IRR, which is inconsistent with the received wisdom.*

We will now discuss a pentad of analyses related to recurring revenue concentration at acquisition and ETA economic outcomes (**Figure 3**).

**Figure 3: The results we will examine**

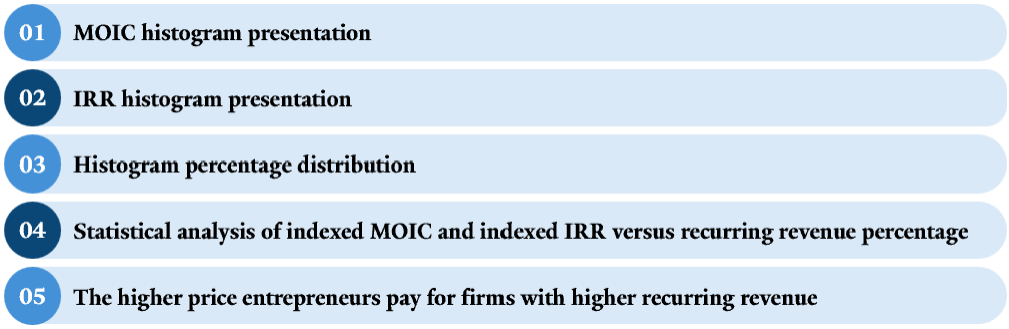


01 *MOIC histogram presentation*

We will begin with a presentation of various histograms to quickly and visually illustrate outcomes by recurring revenue bin. Readers might perceive a positive relationship between recurring revenue percentages and MOICs and IRRs in these histograms. However, the statistical analyses below will reveal that the relationship is weak and inconsistent.

Before sharing our results, we present a fictional bell curve and histogram illustrating what recurring revenue enthusiasts might expect (**Figure 4**). This hypothetical curve and chart show clustering of smaller MOIC outcomes for None and relatively Low recurring revenue categories, and more clustering at robust MOIC levels for relatively Moderate and High recurring revenue categories. Let's see how our hyperbolic depiction aligns with actual observations.

**Figure 4: A recurring revenue aficionado's expected MOIC outcomes**

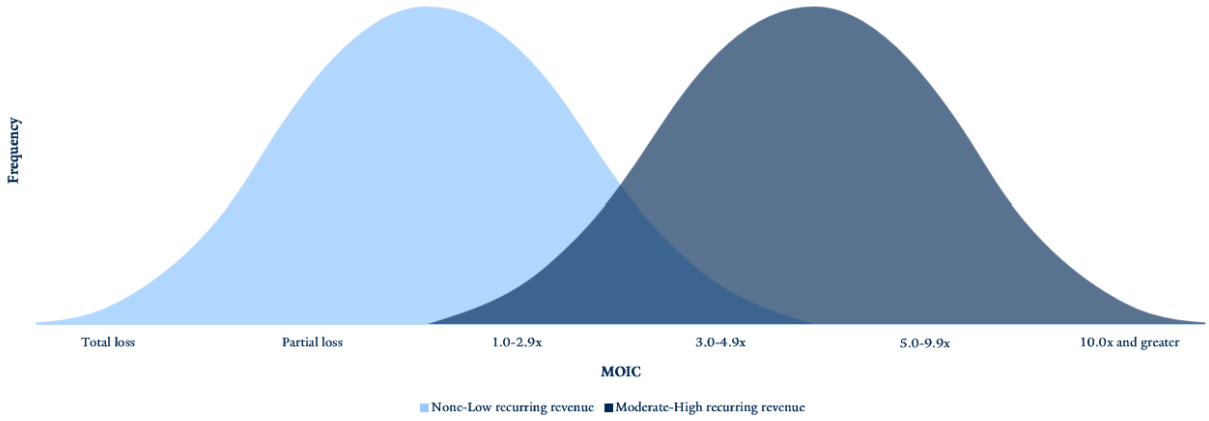

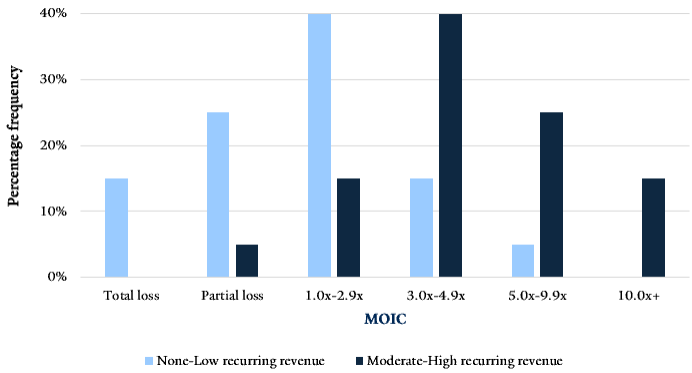


The investor who provided us with the data set granted us permission to report the raw MOIC results by recurring revenue bins (**Figure 5**). Readers can visually compare these actual outcomes to the hypothetical expectations in **Figure 4**. As we move from the High recurring revenue bin (top left) to the None recurring revenue bin (bottom right) in **Figure 5,** there does not seem to be a pronounced leftward shift in the distributions like the one in **Figure 4**.

**Figure 5: Percentage frequencies of MOIC outcomes by recurring revenue bin**

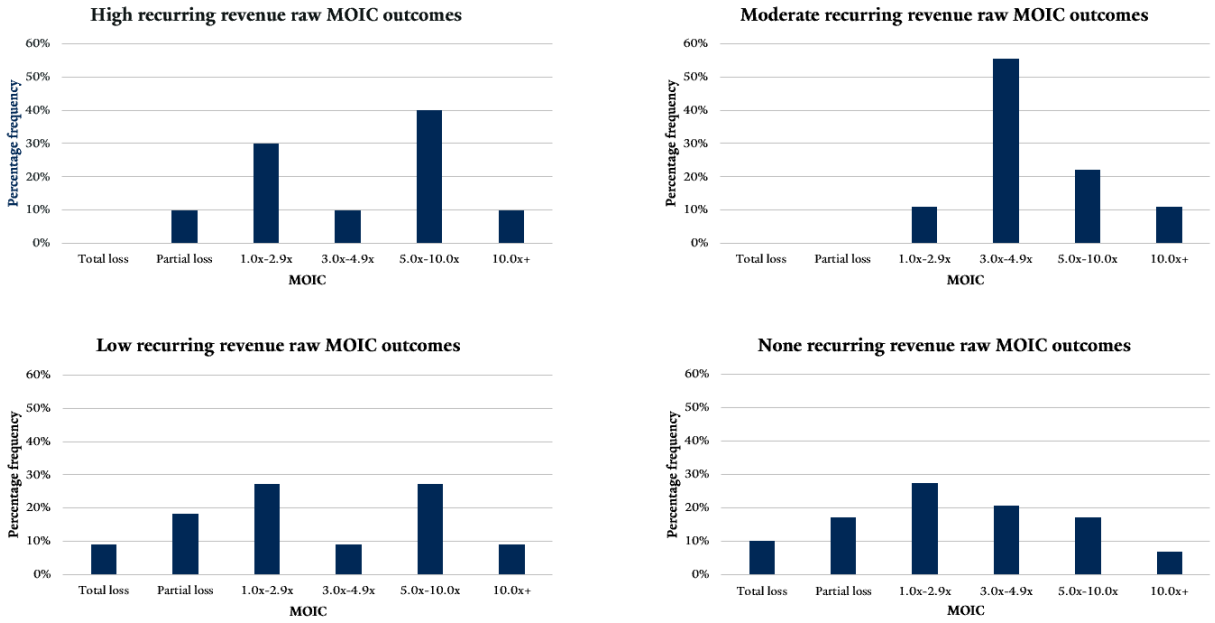


**Figure 6** presents the same data in tabular format. Again, the results are noisy, and it appears that High recurring revenue firms have less deleterious outcomes but not necessarily better upside results. We will test this initial observation below.

**Figure 6: Raw MOIC distribution by recurring revenue category at acquisition (tabular)**

| Raw MOIC distribution | | | | |
|---|---|---|---|---|
| | None | Low | Moderate | High |
| Total loss | 10% | 9% | 0% | 0% |
| Partial loss | 17% | 18% | 0% | 10% |
| 1.0x-2.9x | 28% | 28% | 11% | 30% |
| 3.0x-4.9x | 21% | 9% | 56% | 10% |
| 5.0x-9.9x | 17% | 27% | 22% | 40% |
| 10.0x and higher | 7% | 9% | 11% | 10% |
| Sum | **100%** | **100%** | **100%** | **100%** |

To provide an alternate visual representation, we combined the two bins in the top row of **Figure 5** (High and Moderate recurring revenue) and the two bins in the bottom row (Low and None recurring revenue). In effect, we bifurcated the outcomes at the 50% recurring revenue line. The results reported in **Figure 7** suggest that more recurring revenue at acquisition leads to a fulsome MOIC, somewhat similar to the hypothetical distributions in **Figure 4**. However, as we will demonstrate later, appearances can be deceiving.

**Figure 7: Raw MOIC distribution by recurring revenue category at acquisition (graphical)**

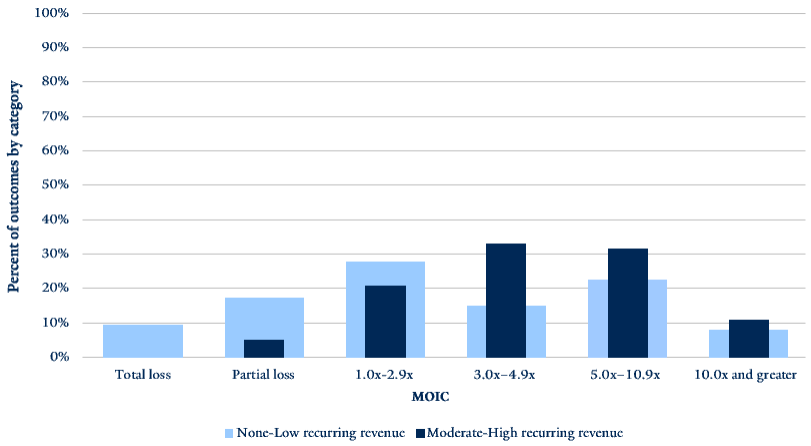


We next display histograms for ten equal-width indexed MOIC groups (**Figure 8**) to supplement the histograms based on raw MOIC buckets (**Figure 5**). In the None recurring revenue category, it looks like 80% of outcomes are poor in the 0–10 indexed range. This is the greatest proportion and compares unfavorably to 50% in the High recurring revenue group. The High and Low recurring revenue batches have no blockbuster outcomes; the Moderate and None recurring revenue sets each have at least one grand slam home run. What does this data tell us so far? Not much.

**Figure 8: Percentage frequency of indexed MOIC results by bin**

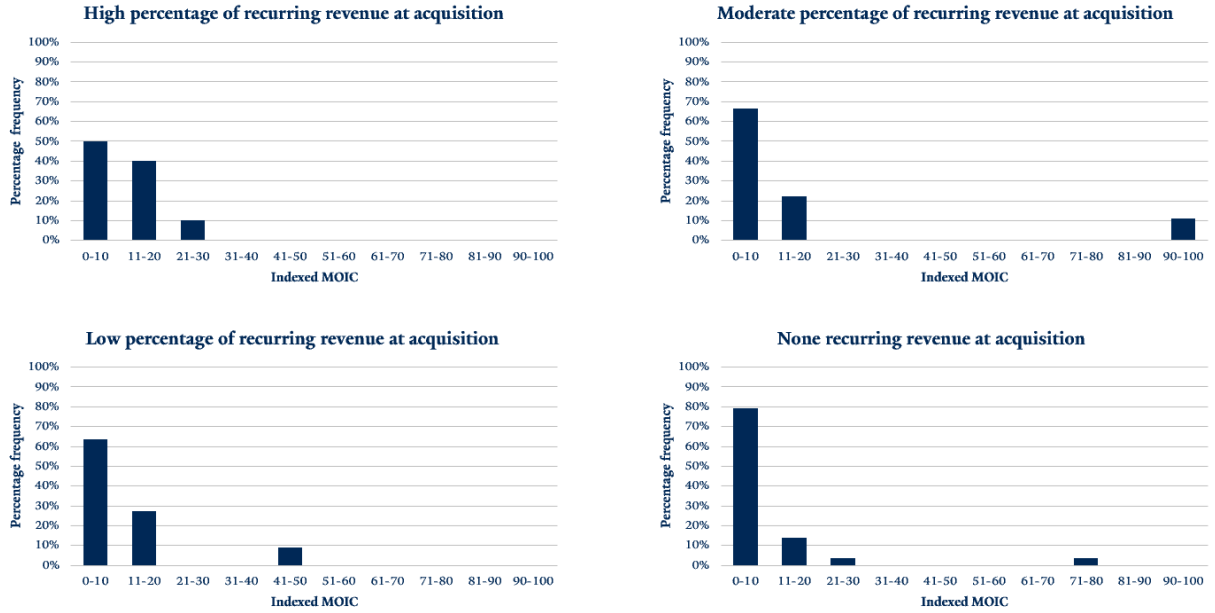


In **Figure 9**, we present the data above using a box plot for an alternative visual display. For readers unfamiliar with box plots, they are a handy tool to present quite a bit of data in an elegant format. They are especially helpful when comparing distributions across categories, as we do here. The top horizontal line indicates the edge of the highest quartile, and the bottom horizontal marks the end of the lowest quartile. The vertical lines (often called whiskers) below and above the blue box show the first (25%) and fourth (75%) quartiles respectively, while the blue box depicts the interquartile range and central tendency. The horizontal line inside the blue box represents the median, and the X mark indicates the mean. With box plots, we can compare medians and variability among different groups. When constructing the box plot, we elected not to include outlier marks as an option in Excel. We did this because the outliers are extreme and distort the visual presentation. The maximum outlier in the None bin has an indexed value of 71; in the Low bin, the zenith outlier has an indexed value of 47, and in the Moderate bin, the peak outlier lies at an indexed value of 100. The High bin has no outlier values.

There appears to be a general upward trend in indexed MOIC as recurring revenue amplifies. However, two caveats exist. First, the upward trend seems to reverse after the Moderate bin. Second, indexed MOIC values are quite spread out within each bin (as shown in **Figures 5** and **6**). This suggests that the overall upward trend might be insignificant when using statistical methods that account for this dispersion in outcomes.

**Figure 9: Indexed MOIC analysis in box plot format**

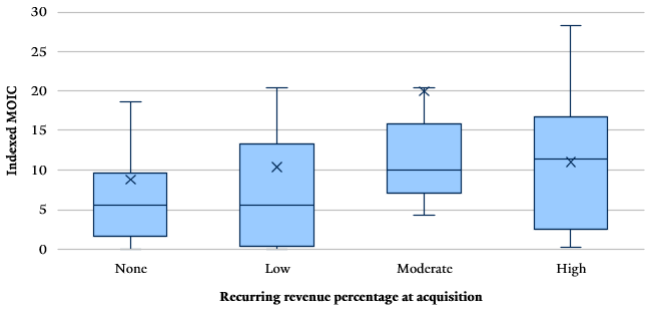


02 *IRR histogram presentation*

For the IRR analysis, we will begin with percentage frequency distributions across revenue bins (**Figure 10**), as we did with the MOIC analysis. Visually, it appears that the Moderate recurring revenue bin, not the High bin, tilts toward upper IRR outcomes. We will test this below.

**Figure 10: Percentage frequencies of IRR outcomes by recurring revenue bin**

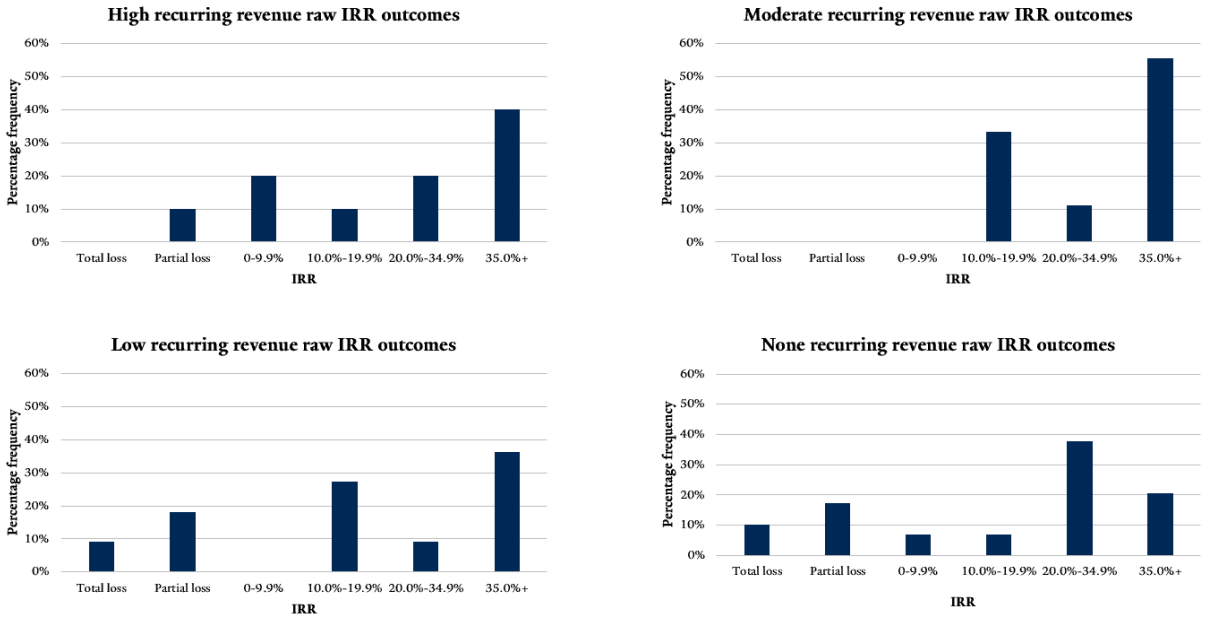


Similar to **Figure 7**, in **Figure 11**, we combine the None and Low revenue bins and the Moderate and High revenue bins in a graphical depiction of IRR outcomes. While it appears that the None and Low bins experience more total and partial losses and that the Moderate and High bins enjoy more favorable IRR outcomes, there remains considerable noise and randomness, as we will demonstrate below using statistical tools.

**Figure 11: Percentage IRR distribution by recurring revenue category at acquisition**

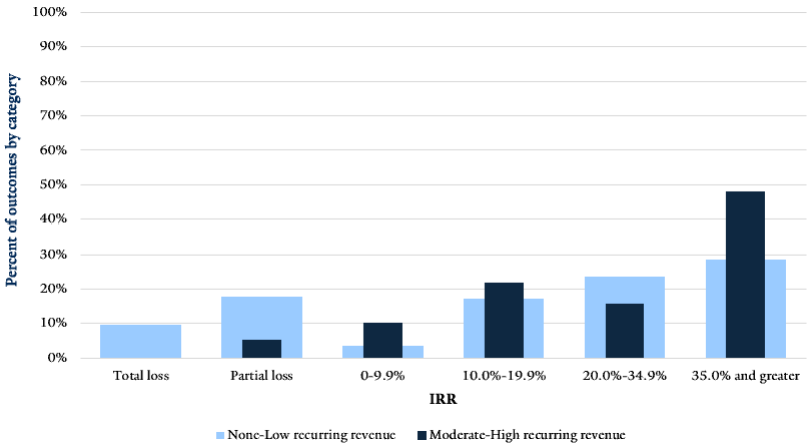


We analyzed histograms of indexed IRR results, expecting to see patterns similar to the indexed MOIC plots in **Figure 8**. Generally, this is true. However, indexed IRR offers a slightly different perspective: If a company has a short holding period with a reasonable MOIC, the IRR might appear towering. Conversely, if a company has a long holding period with a strong MOIC, the IRR might seem relatively mundane. In **Figure 12**, we provide a different perspective by showing histograms based on ten equal-width indexed IRR groups, rather than the raw IRR buckets.

**Figure 12: Percentage frequency of indexed IRR results by bin**

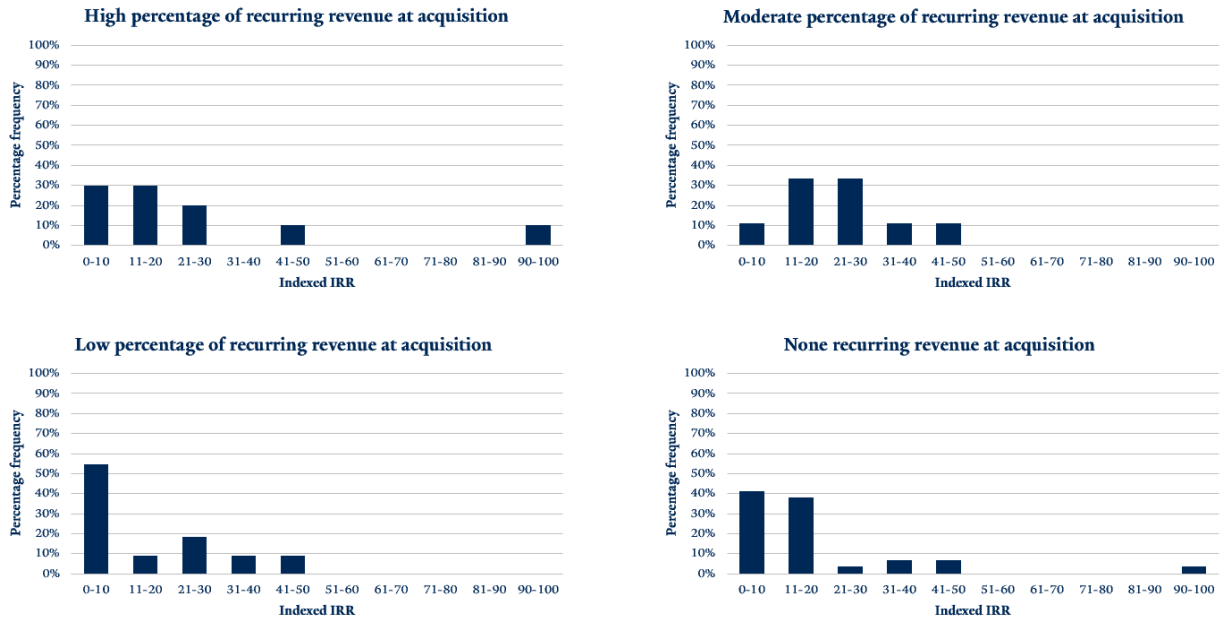


Directionally, this mirrors the indexed MOIC plots in **Figure 8** above. However, there seems to be a significant indexed IRR event in the High recurring revenue bracket that did not show up in the indexed MOIC data. This suggests a shorter holding period. Similarly, the Moderate array lost an outlier in the indexed IRR plot that was present in the indexed MOIC histogram, indicating a longer hold period. Finally,

the None recurring revenue panel displays a single outlier winner as seen in the indexed MOIC display. Once again, what does this data tell us? Not much.

**Figure 13** displays the indexed IRR data in a box plot format. Similar to the MOIC plots in **Figure 9**, the indexed IRR plots show a general upward trend as recurring revenue increases, but the trend is again associated with substantial uncertainty, as indicated by the spread in IRR for all recurring revenue bins. A closer look reveals some differences compared to **Figure 9**: The median indexed IRR exhibits an erratic pattern, declining between None and Low recurring revenue, and again between Moderate and High recurring revenue. Like **Figure 9**, we excluded outlier values for a better visual presentation. The None and High bins each had one outlier observation with an indexed value of 100.

**Figure 13: Indexed IRR analysis in box plot format**

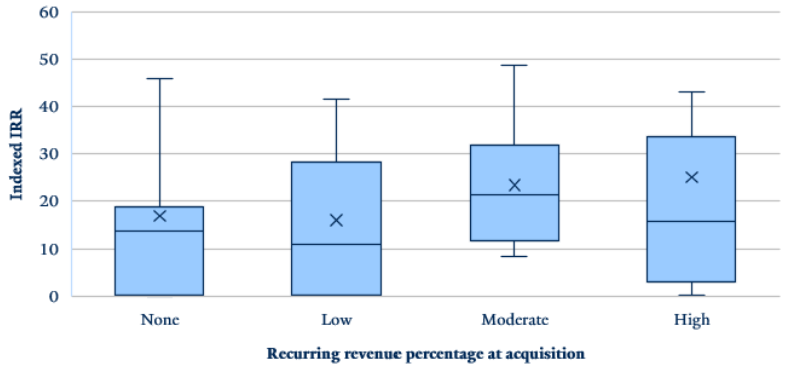


03 *Histogram percentage distribution*

We present the relative percentage histogram distributions from **Figures 8** and **12** as tables in **Figure 14**. These indexed MOIC and IRR tables offer a quick and simple way for readers to see distribution density in a tabular format. Looking at the left panel, readers might jump to the conclusion that firms in the None recurring revenue bin have weaker outcomes. After all, 79% of the observations are in the 0–10 bin, while only 50% of the High recurring revenue outcomes fall into the same tier. However, if we combine the bottom two levels to create a 0–20 class, the High cluster sums to 90%, while the None group is at 93%, making them nearly indistinguishable. Additionally, keep in mind that entrepreneurs who generate low raw MOICs, such as 0–2, do not fare well in the ETA game. They might earn a few bucks, but the result is modest and not the aspirational outcome they seek. Similarly, investors earn some proceeds from low raw MOIC outcomes, but such outcomes dilute overall fund returns and are not the winners.

Let's shift to the big indexed MOIC outcomes. It appears that the High category has none, the Moderate category has one, the Low category has none, and the None pack has one. Not very revealing data. We will skip the right-indexed IRR panel, which shows a similar distribution to the indexed MOIC.

So, what can we learn from this? We think readers might mistakenly jump to the conclusion that High recurring revenue enterprises have fewer negative observations and that there is optionality on the upside, so why not focus on recurring revenue? Consistent with our discussion of the results in **Figures 9** and **13**, the distributions in **Figure 14** reinforce that view. However, that conclusion is premature, as we will shortly demonstrate with statistical analysis.

**Figure 14: Relative percentage distributions for indexed MOIC and indexed IRR outcomes**

| Indexed MOIC distribution | | | | |
|---|---|---|---|---|
| | None | Low | Moderate | High |
| 0-10 | 79% | 60% | 70% | 50% |
| 11-20 | 15% | 30% | 20% | 40% |
| 21-30 | 3% | 0% | 0% | 10% |
| 31-40 | 0% | 0% | 0% | 0% |
| 41-50 | 0% | 10% | 0% | 0% |
| 51-60 | 0% | 0% | 0% | 0% |
| 61-70 | 0% | 0% | 0% | 0% |
| 71-80 | 3% | 0% | 0% | 0% |
| 81-90 | 0% | 0% | 0% | 0% |
| 90-100 | 0% | 0% | 10% | 0% |
| Sum | 100% | 100% | 100% | 100% |

| Indexed IRR distribution | | | | |
|---|---|---|---|---|
| | None | Low | Moderate | High |
| 0-10 | 41% | 50% | 20% | 30% |
| 11-20 | 39% | 10% | 30% | 30% |
| 21-30 | 3% | 20% | 30% | 20% |
| 31-40 | 7% | 10% | 10% | 0% |
| 41-50 | 7% | 10% | 10% | 10% |
| 51-60 | 0% | 0% | 0% | 0% |
| 61-70 | 0% | 0% | 0% | 0% |
| 71-80 | 0% | 0% | 0% | 0% |
| 81-90 | 0% | 0% | 0% | 0% |
| 90-100 | 3% | 0% | 0% | 10% |
| Sum | 100% | 100% | 100% | 100% |

04 *Statistical analysis of indexed MOIC and indexed IRR versus recurring revenue percentage*

We will now shift focus to discuss the statistical analyses conducted on our data set. **Figure 15** plots the relationship between indexed MOIC and recurring revenue percentage at acquisition. It underpins the correlation and regression results reported in **Figure 17** below. Supporters of recurring revenue companies would expect to see a positive relationship: more points in the lower left and upper right quadrants of the plot, showing that when recurring revenue is low (high), indexed MOIC is also low (high). But that is not the case. Although the fitted (regression) line slope is slightly positive at 6.4 in **Figure 17**, the t-stat of 1.10 and p-value of 0.27 indicate there is no reliable statistical relationship here.

**Figure 15: Indexed MOIC versus recurring revenue plot**

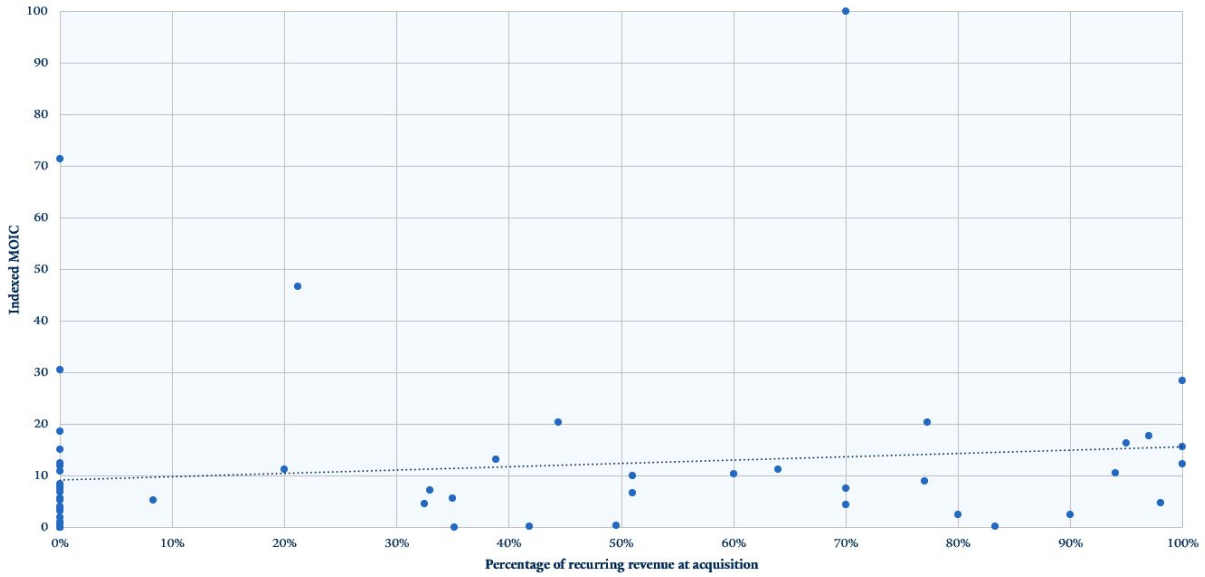


We observe the same pattern in **Figure 16**, which plots indexed IRR against recurring revenue at acquisition. There is a positive slope (9.48), but the low t-stat of 1.32 and high p-value of 0.19 reveal that the slope is not significant. The results seem random. It is true that there is a notable outlier in the upper

right corner—a firm with 100% recurring revenue that achieved an indexed IRR of 100. However, there is also a firm with 0% recurring revenue, appearing in the upper left corner with the same indexed IRR.

**Figure 16: Indexed IRR versus recurring revenue plot**

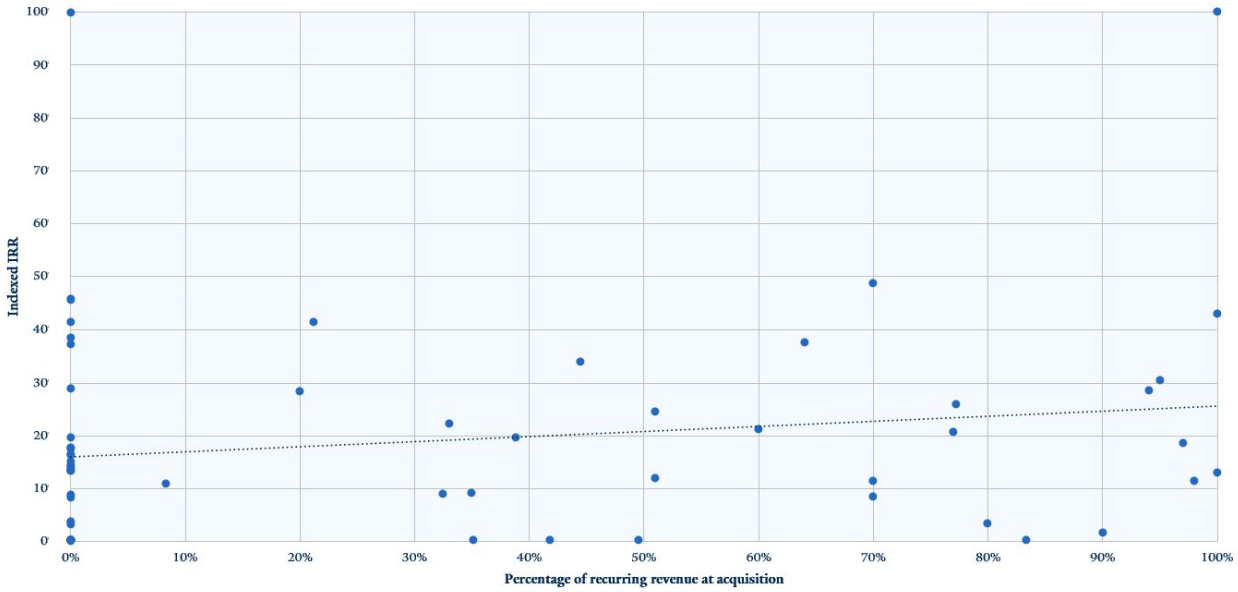


To address potential concerns that our results for regressions and differences in mean outcomes across recurring revenue bins are influenced by outliers, we repeated the analyses based on ranks of MOIC, IRR, and recurring revenue. We find similar results, indicating that our inability to identify a significant relationship between financial outcomes and recurring revenue is not due to outliers. For those curious about the ranked plots, they can be found in **Exhibit 2**.

**Figure 17** summarizes key statistical points for the overall distributions of indexed MOIC and IRR in the first and third columns. We marked the data with two darker horizontal lines. The numbers above the first line relate solely to the Y variables. The data below the first line pertains to the relationship between the Y variables and the percentage of recurring revenue at acquisition, including the paired correlation and p-value between the two lines. Both means are higher than the respective medians, indicating right-skewed distributions. Regarding the relationships between outcomes and recurring revenue, we report correlations between indexed MOIC and IRR and recurring revenue percentages, as well as the slopes from regressions of indexed MOIC and IRR on recurring revenue percentages. The correlations are positive but not statistically significant. The regression line slopes are also positive but again insignificant: The t-statistics are low, and p-values are all high. For example, in the indexed IRR column, the high p-value of 0.19 associated with the correlation of 0.17 indicates that there is no statistical relationship. The results in the second and fourth columns describe a similar analysis conducted on ranked MOIC and IRR, which is designed to reduce the impact of outliers. Even though the plots in **Exhibit 2** suggest a positive relation between recurring revenue and ranked MOIC and IRR, our results confirm that the relation is not statistically significant, indicated by relatively low t-statistics on the slopes. Therefore, *we cannot reject the null hypothesis that there is no relationship between recurring revenue concentration and MOIC and IRR results for ETA firms.*

**Figure 17: Summary statistical data**

| | Indexed MOIC | Ranked MOIC | Indexed IRR | Ranked IRR |
|---|---|---|---|---|
| Mean | 11.2 | 30.0 | 19.11 | 30.0 |
| Median | 7.3 | 30.0 | 14.14 | 30.0 |
| Standard deviation | 16.6 | 17.2 | 20.47 | 17.2 |
| Interquartile range | 8.8 | 29.0 | 23.61 | 0.4 |
| Correlation | 0.14 | 0.26 | 0.17 | 0.19 |
| p-value | 0.27 | 0.04 | 0.19 | 0.14 |
| Slope | 6.45 | 12.09 | 9.48 | 8.94 |
| Standard error | 5.84 | 5.88 | 7.16 | 5.97 |
| t-test | 1.10 | 2.05 | 1.32 | 1.49 |
| Rsq | 0.02 | 0.06 | 0.02 | 0.03 |
| Count | 59 | 59 | 59 | 59 |

The statistical analyses we have discussed so far incorporate the recurring revenue properties of individual firms, rather than groupings based on the four recurring revenue bins. We next examine whether there are statistically significant outcome differences across these bins. A summary of the key statistics is provided in **Figure 18**. Note that the means, medians, and interquartile ranges have already been shown in the box plots in **Figures 9** and **13**. Once again, readers might be tempted to conclude that financial outcomes improve with recurring revenue. For example, the High and Moderate recurring revenue categories have more advanced average and median indexed MOICs than the Low and None categories. A similar visual pattern appears for indexed IRR results. The real question is whether these differences are statistically significant or just random noise. We will discuss that now.

**Figure 18: Summary of statistical data by recurring revenue bin**

| | Mean Indexed MOIC | Median Indexed MOIC | SD Indexed MOIC | Interquartile Indexed MOIC | Mean Indexed IRR | Median Indexed IRR | SD Indexed IRR | Interquartile Indexed IRR |
|---|---|---|---|---|---|---|---|---|
| High | 11.1 | 11.4 | 8.8 | 13.0 | 25.1 | 15.8 | 29.8 | 24.7 |
| Moderate | 19.9 | 10.1 | 30.4 | 3.6 | 23.4 | 21.3 | 13.1 | 13.9 |
| Low | 10.4 | 5.6 | 13.5 | 9.7 | 16.0 | 10.9 | 14.3 | 20.6 |
| None | 8.8 | 5.6 | 13.7 | 6.4 | 16.9 | 13.7 | 20.8 | 17.5 |

While the regression results consistently show no significant link between value creation and recurring revenue, readers might wonder if performance differs significantly between the two extreme recurring revenue subgroups: None recurring revenue and High recurring revenue. Once again, the results in **Figures 19** and **20** indicate no significant relationship. We first performed a Mood's Median Test to see if the median indexed MOIC for the High recurring revenue bin is greater than the median for the None recurring revenue bin. For those unfamiliar with Mood's test, it is a statistical method used to determine if there is any significance in the relationship between the medians of two samples. To perform the test, we pooled all indexed MOIC observations from both the High and None recurring revenue groups and calculated the median for the combined data. Then, we counted how many observations in each subgroup were above and below the median. We calculated the expected values and conducted a chi-squared test to determine a p-value. We used the same process for indexed IRR observations. The results are shown in **Figure 19**. With p-values of 0.40 and 0.92 for indexed MOIC and indexed IRR respectively, we find no significant difference between the medians of the High and None recurring revenue groups.

**Figure 19: Mood's Median Test results from comparing median outcomes in High and None recurring bins**

**Indexed MOIC test**

| | Observed Values | | | Expected Values | | |
|---|---|---|---|---|---|---|
| | Above | Below | Sum | Above | Below | Sum |
| High | 6.00 | 4.00 | 10.00 | 4.87 | 5.13 | 10.00 |
| None | 13.00 | 16.00 | 29.00 | 14.13 | 14.87 | 29.00 |
| **Sum** | **19.00** | **20.00** | **39.00** | **19.00** | **20.00** | **39.00** |
| Indexed MOIC median | 5.67 | | | | | |
| Chi-squared | 0.68 | | | | | |
| p-value | 0.40 | | | | | |

**Indexed IRR test**

| | Observed Values | | | Expected Values | | |
|---|---|---|---|---|---|---|
| | Above | Below | Sum | Above | Below | Sum |
| High | 5.00 | 5.00 | 10.00 | 4.87 | 5.13 | 10.00 |
| None | 14.00 | 15.00 | 29.00 | 14.13 | 14.87 | 29.00 |
| **Sum** | **19.00** | **20.00** | **39.00** | **19.00** | **20.00** | **39.00** |
| Indexed IRR median | 0.14 | | | | | |
| Chi-squared | 0.00 | | | | | |
| p-value | 0.92 | | | | | |

We also performed a t-test to compare the means for the High recurring revenue and None recurring revenue bins (**Figure 20**). The t-stat is below the t critical values in all cases, producing high p-values, which indicates no significant difference between the means of the two bins. *In other words, we cannot reject the null hypothesis that the MOICs and IRRs are the same between the High and None recurring revenue clusters*.

**Figure 20: t-test results comparing mean outcomes in High and None recurring revenue bins**

| | Indexed MOIC | | Indexed IRR | |
|---|---|---|---|---|
| | High | None | High | None |
| Mean | 11.09 | 8.84 | 0.25 | 0.16 |
| Standard deviation | 8.81 | 13.73 | 0.29 | 0.20 |
| Observations | 10 | 29 | 10 | 29 |
| Hypothesized mean difference | 0 | | 0 | |
| Degrees of freedom | 25 | | 12 | |
| t Stat | 0.59 | | 0.80 | |
| P(T<=t) one-tail | 0.27 | | 0.21 | |
| t Critical one-tail | 1.70 | | 1.78 | |
| P(T<=t) two-tail | 0.55 | | 0.43 | |
| t Critical two-tail | 2.05 | | 2.17 | |

In all of our analyses, there is no statistically significant linear relationship between the percentage of recurring revenue and MOIC or IRR outcomes. To be sure, readers might see a general positive relation between MOIC, IRR, and recurring revenue, or even notice a nonlinear relationship, with the sweet spot lying in the Moderate range (50% to 80% recurring revenue). However, our statistical analyses indicate there is substantial randomness underlying these patterns, making them unreliable for entrepreneurs and investors.

Although we believe our data set is extensive, we wish we had more observations. All of our analyses currently suggest weak coefficients and high p-values. With more data and additional explanatory variables, our conclusions would likely become less tentative. We might continue to see a weak relationship between recurring revenue and economic outcomes if the coefficient statistics remain small and the p-values stay

high, supporting the current conclusion. Conversely, more data could uncover a genuine positive relationship, reducing standard errors and pushing the coefficients toward significance, indicating that recurring revenue does matter. Alternatively, additional data could also reveal a true negative relationship, showing a significant inverse effect.

Whatever the case, we hope our data offers a note of caution for ETA entrepreneurs and investors. Maybe recurring revenue is not a magic solution or a reliable indicator of future returns.

05 *The higher price entrepreneurs pay for firms with higher recurring revenue*

In all of our analyses, we found only one relationship that was statistically significant—the percentage of recurring revenue at acquisition and entry EBITDA multiples. Entrepreneurs pay more for target acquisitions with additional percentages of recurring revenue. In our analysis, we eliminated observations with no positive EBITDA dollars at acquisition; we do not wish to punish companies that might trade on a revenue basis and do not have earnings. **Figure 21** illustrates this very positive relationship with a regression line at 1.61, a correlation of 0.36, a t-stat of 2.83, and a low p-value of 0.006. This gives us more than 99% confidence that entry EBITDA multiples increase with recurring revenue. The takeaway is that part of the reason why High recurring revenue firms might not deliver superior MOICs and IRRs is that entrepreneurs are paying more to acquire those assets, creating a steeper hurdle to achieving greater absolute returns. **Figure 22** displays the data in a box plot. The Moderate bin has an outlier observation with an EBITDA entry multiple that just exceeds 10.

**Figure 21: Entry EBITDAx versus recurring revenue percentage at acquisition (with no zero EBITDA observations)**

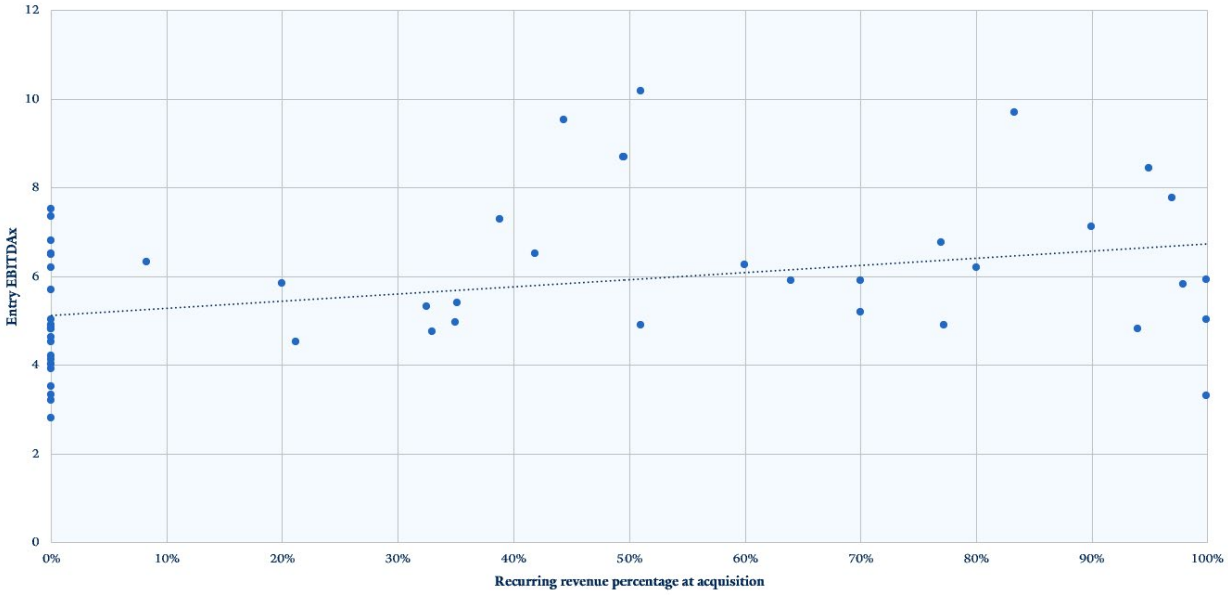

**Figure 22: Entry EBITDAx versus recurring revenue percentage at acquisition (with no zero EBITDA observations)**

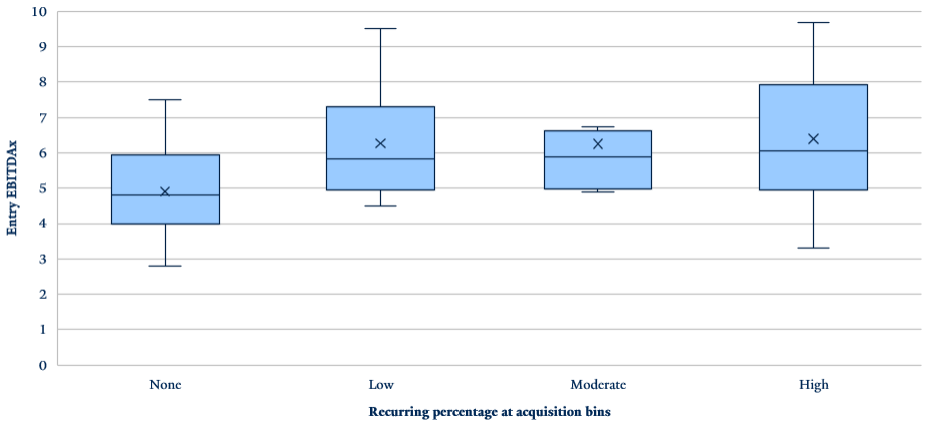


Additionally, we examined the relationship between the percentage of recurring revenue at acquisition and IRR, while controlling for the entry EBITDA multiple, to assess whether a clear connection exists between these outcomes and recurring revenue. The result is: *no*. The IRR regression coefficient flips to a negative value (-0.12) but remains insignificant with a very high p-value of 0.44, reinforcing the notion that the percentage of recurring revenue at acquisition has no statistically significant association with IRRs.

## Why recurring revenue might still be important

Our data suggests that recurring revenue is not the miracle cure that most ETA entrepreneurs and investors believe it to be. However, this does not mean recurring revenue should be avoided or overlooked. We cannot make that claim. Recurring revenue in search fund projects can still be highly valuable, and we believe there are at least three reasons to consider embracing recurring revenue (**Figure 23**).

**Figure 23: A triad of reasons to still pursue recurring revenue despite no apparent correlation with MOIC and IRR**

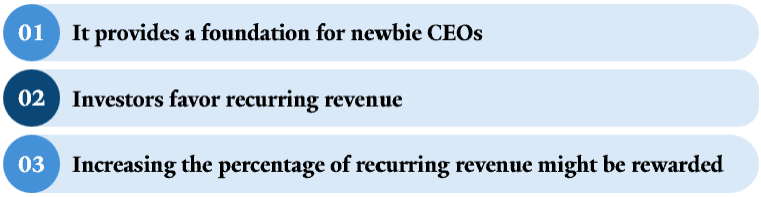


01 *It provides a foundation for newbie CEOs*

When fresh CEOs take the helm of a small, squirrely company, they can feel overwhelmed by managing operations, understanding the business, getting to know employees, and building relationships with existing customers. It is a lot to manage. One key advantage of a business with robust recurring revenue is that it has an established customer base providing a steady stream of top-line income, which can ideally lead to healthy bottom-line cash flow. The benefit of this setup is that it allows the rookie CEO to find their footing and get used to the CEO role, which is a significant plus since most ETA entrepreneurs have never been CEOs before.

The critical advantage of a recurring revenue model is that the CEO does not have to wake up each day worried about and scrambling for revenue and customers; they are already in place. In a transactional revenue business, each day starts without guaranteed sales, and the CEO and team must focus on acquiring new customers. While this might not negatively impact MOICs and IRRs, it changes how inexperienced CEOs spend their time. In a recurring revenue enterprise, CEOs do not need to prioritize sales generation in the earliest innings. This is a wonderful luxury for a neophyte leader.

02 *Investors favor recurring revenue*

Although we cannot provide an exact number of investors in the ETA community, we estimate there are hundreds. Attend any search fund conference, and the halls are packed with smart capital providers selling money and wooing potential entrepreneurs. This group is at the ready and eager to finance entrepreneurs who meet their profile criteria, as well as company acquisitions that fit within their target deal scope. The system is awash with funds for the right people and the right deals. Money is not the limiting factor; a top-notch target opportunity at a fair trading price is always the bottleneck. In other words, funding for ETA acquisitions is a pretty thick market; for the right deal, money is readily available.

One of the criteria that ETA investors perceive as meeting their profile is recurring revenue. Whether or not this is logical regarding economic outcomes, it is what capital providers crave. Entrepreneurs face a steep climb in the ETA trek, and one way to make it easier is to pitch poster child deals to investors. If they seek recurring revenue, it simplifies the process, and the deal is more likely to get funded. In the current ETA funding market, entrepreneurs are incentivized to find recurring revenue prospects to purchase.

03 *Increasing the percentage of recurring revenue might be rewarded*

Our data set only measures the percentage of recurring revenue in a target acquisition at the inception of the ETA CEO's tenure. We lack information on how this percentage *changed* throughout the investment's holding period. We suspect that firms increasing their recurring revenue percentage might be rewarded by the next buyer, as the business is reimagined with a characteristic that typically commands topping EBITDA multiples, as our data demonstrates.

This suggests that if an ETA CEO can transition a business from Moderate recurring revenue to High recurring revenue, there is potential for financial gain. We often hear about software companies moving from a license revenue model (a one-time fee) to a monthly or yearly subscription model. While this is a very challenging bridge to cross because it causes a short-term cash shortfall, it can be quite lucrative if successful.

We have never claimed or implied that recurring revenue is harmful. On the contrary, we are fans—especially if the CEO can augment the percentage during their tenure. That boosts the company's value, which is often rewarded. This differs from buying a business with an elevated level of recurring revenue (and paying a premium for it). In that case, there is limited room for improvement.

## Conclusion

*Despite high percentage recurring revenue firms being popular acquisition targets in the ETA pond, they do not link to proportionately superior economic outcomes, measured by MOIC and IRR, compared to firms with low or no recurring revenue*. Our data and analyses (histograms, regressions, Mood's tests, and t-tests) reveal there is no reliable relationship between recurring revenue percentages at acquisition and financial outcomes. This challenges a core tenet in the search fund community.

Regardless of our findings, there are still valid reasons for ETA entrepreneurs to chase recurring revenue opportunities. They provide a stable foundation for transitioning into the CEO role, they are attractive to investors, and there is an opportunity to create value by increasing the proportion of recurring revenue. However, entrepreneurs should be prepared to pay lofty EBITDA multiples for the chance to own a substantial percentage recurring revenue business.

We hope our note is at least somewhat provocative and encourages ETA entrepreneurs and investors to rethink the prevailing conventional wisdom surrounding recurring revenue. It is possible that relaxing the recurring revenue constraint could lead to more deals closing or to lower initial investment multiples. We believe entrepreneurs can broaden their view of potential industries without risking their chances of becoming wealthy.

No matter where you stand on recurring revenue, we wish you fun, fulfillment, and MOICs and IRRs that are stratospheric. We hope good luck and good fortune are headed your way, and as always, we are cheering you on every step of the journey!

## Exhibit 1: Our data set

### High percentage of recurring revenue at acquisition

| Recurring revenue % | Indexed MOIC | Indexed IRR | Ranked recurring revenue | Ranked MOIC | Ranked IRR |
|---|---|---|---|---|---|
| 100% | 28 | 43 | 58 | 55 | 55 |
| 100% | 16 | 100 | 58 | 49 | 59 |
| 100% | 12 | 13 | 58 | 45 | 26 |
| 98% | 5 | 11 | 56 | 22 | 23 |
| 97% | 18 | 19 | 55 | 51 | 37 |
| 95% | 16 | 31 | 54 | 50 | 48 |
| 94% | 10 | 29 | 53 | 40 | 46 |
| 90% | 2 | 2 | 52 | 13 | 13 |
| 83% | 0 | 0 | 51 | 8 | 5 |
| 80% | 3 | 3 | 50 | 14 | 15 |

### Moderate percentage of recurring revenue at acquisition

| Recurring revenue % | Indexed MOIC | Indexed IRR | Ranked recurring revenue | Ranked MOIC | Ranked IRR |
|---|---|---|---|---|---|
| 77% | 20 | 26 | 49 | 53 | 44 |
| 77% | 9 | 21 | 48 | 37 | 40 |
| 70% | 100 | 49 | 46 | 59 | 57 |
| 70% | 8 | 8 | 46 | 31 | 18 |
| 70% | 4 | 11 | 46 | 20 | 24 |
| 64% | 11 | 38 | 44 | 42 | 51 |
| 60% | 10 | 21 | 43 | 39 | 41 |
| 51% | 10 | 25 | 42 | 38 | 43 |
| 51% | 7 | 12 | 42 | 28 | 25 |

### Low percentage of recurring revenue at acquisition

| Recurring revenue % | Indexed MOIC | Indexed IRR | Ranked recurring revenue | Ranked MOIC | Ranked IRR |
|---|---|---|---|---|---|
| 49% | 0 | 0 | 40 | 9 | 10 |
| 44% | 20 | 34 | 39 | 54 | 49 |
| 42% | 0 | 0 | 38 | 7 | 10 |
| 39% | 13 | 20 | 37 | 47 | 39 |
| 35% | 0 | 0 | 36 | 2.5 | 5 |
| 35% | 6 | 9 | 35 | 25 | 21 |
| 33% | 7 | 22 | 34 | 30 | 42 |
| 33% | 5 | 9 | 33 | 21 | 20 |
| 21% | 47 | 42 | 32 | 57 | 54 |
| 20% | 11 | 28 | 31 | 43 | 45 |
| 8% | 5 | 11 | 30 | 24 | 22 |

### None recurring revenue at acquisition

| Recurring revenue % | Indexed MOIC | Indexed IRR | Ranked recurring revenue | Ranked MOIC | Ranked IRR |
|---|---|---|---|---|---|
| 0% | 71 | 39 | 15 | 58 | 52 |
| 0% | 30 | 15 | 15 | 56 | 32 |
| 0% | 19 | 42 | 15 | 52 | 53 |
| 0% | 15 | 46 | 15 | 48 | 56 |
| 0% | 12 | 29 | 15 | 46 | 47 |
| 0% | 12 | 13 | 15 | 44 | 28 |
| 0% | 11 | 17 | 15 | 41 | 34 |
| 0% | 8 | 16 | 15 | 36 | 33 |
| 0% | 8 | 18 | 15 | 35 | 36 |
| 0% | 8 | 100 | 15 | 34 | 58 |
| 0% | 8 | 37 | 15 | 33 | 50 |
| 0% | 8 | 14 | 15 | 32 | 30 |
| 0% | 7 | 14 | 15 | 29 | 29 |
| 0% | 6 | 13 | 15 | 27 | 27 |
| 0% | 6 | 17 | 15 | 26 | 35 |
| 0% | 5 | 20 | 15 | 23 | 38 |
| 0% | 4 | 14 | 15 | 19 | 31 |
| 0% | 4 | 4 | 15 | 18 | 16 |
| 0% | 4 | 3 | 15 | 17 | 14 |
| 0% | 3 | 8 | 15 | 16 | 17 |
| 0% | 3 | 9 | 15 | 15 | 19 |
| 0% | 2 | 0 | 15 | 12 | 1 |
| 0% | 1 | 0 | 15 | 11 | 10 |
| 0% | 1 | 0 | 15 | 10 | 10 |
| 0% | 0 | 0 | 15 | 6 | 5 |
| 0% | 0 | 0 | 15 | 5 | 10 |
| 0% | 0 | 0 | 15 | 2 | 5 |
| 0% | 0 | 0 | 15 | 2 | 5 |
| 0% | 0 | 0 | 15 | 2 | 5 |

**Exhibit 2: Ranked analyses of MOIC and IRR versus recurring revenue**

Plot A: Ranked MOIC versus ranked recurring revenue plot

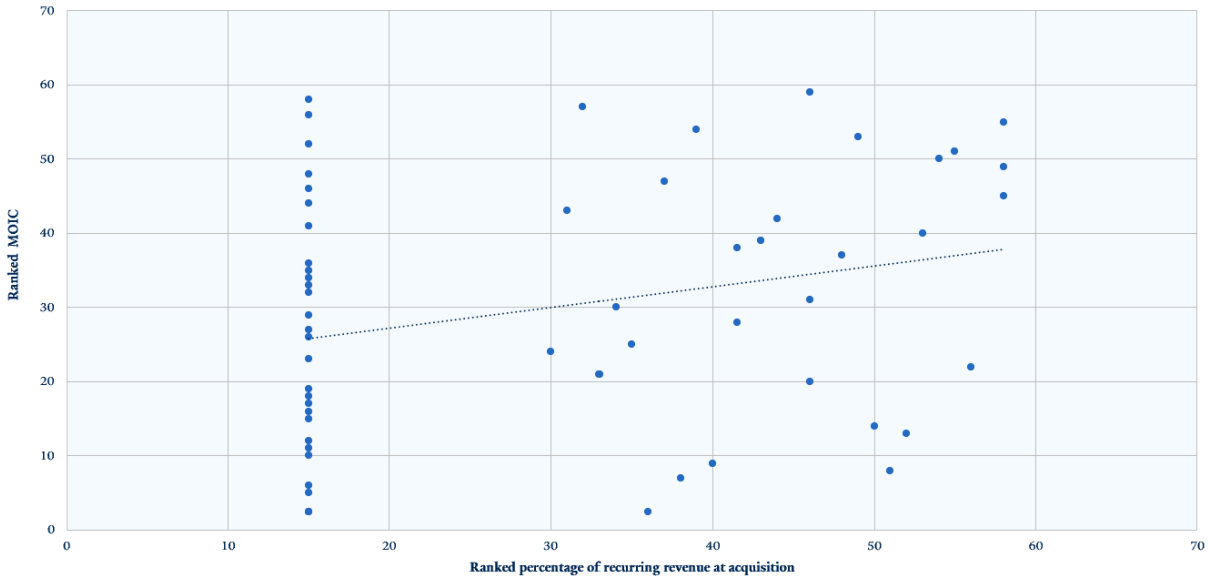


Plot B: Ranked IRR versus ranked recurring revenue plot

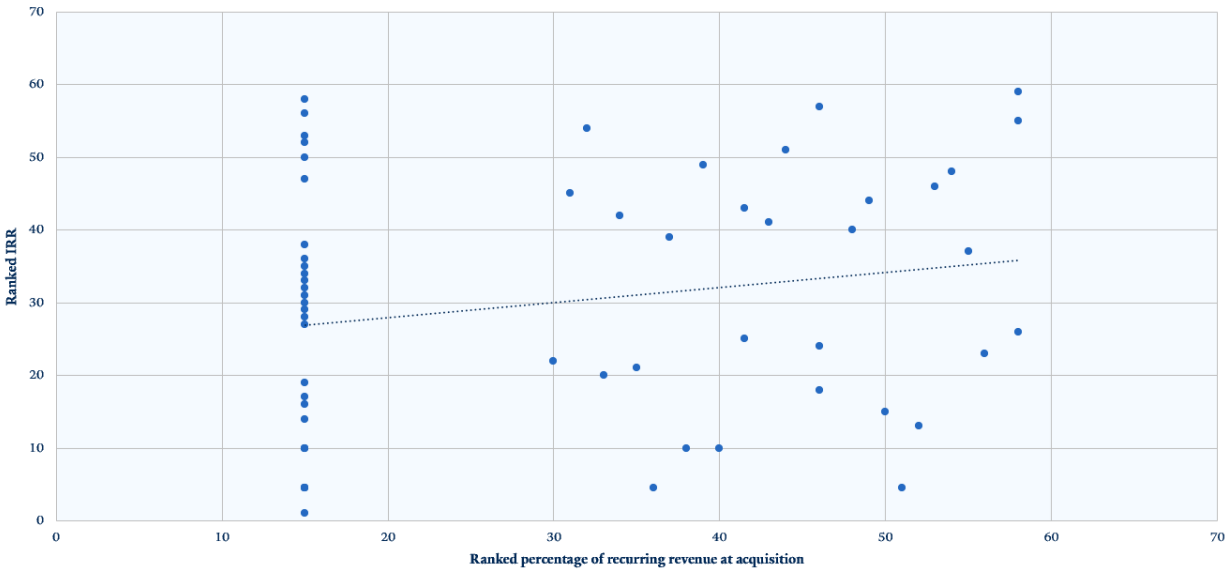

## Exhibit 3: A brief glossary of statistical terms

We directly reproduced definitions from stat.berkeley.edu/~stark/SticiGui/Text/gloss.htm. The definitions below are not our original work and are attributable to Berkeley's very helpful statistics glossary page.

*Correlation*: A measure of linear association between two (ordered) lists. Two variables can be strongly correlated without having any causal relationship, and two variables can have a causal relationship and yet be uncorrelated.

*Histogram*: A histogram is a kind of plot that summarizes how data are distributed. Starting with a set of class intervals, the histogram is a set of rectangles ("bins") sitting on the horizontal axis. The bases of the rectangles are the class intervals, and their heights are such that their areas are proportional to the fraction of observations in the corresponding class intervals. That is, the height of a given rectangle is the fraction of observations in the corresponding class interval, divided by the length of the corresponding class interval. A histogram does not need a vertical scale, because the total area of the histogram must equal 100%. The units of the vertical axis are percent per unit of the horizontal axis. This is called the density scale. The horizontal axis of a histogram needs a scale. If any observations coincide with the endpoints of class intervals, the endpoint convention is important.

*Interquartile range*: The interquartile range of a list of numbers is the upper quartile minus the lower quartile.

*Mean*: The sum of a list of numbers, divided by the number of elements in the list.

*Median*: The middle value of a list. The smallest number such that at least half the numbers in the list are no greater than it. If the list has an odd number of entries, the median is the middle entry in the list after sorting the list into increasing order. If the list has an even number of entries, the median is the smaller of the two middle numbers after sorting. The median can be estimated from a histogram by finding the smallest number such that the area under the histogram to the left of that number is 50%

*p-value*: Suppose we have a family of hypothesis tests of a null hypothesis that let us test the hypothesis at any significance level p between 0 and 100% we choose. The p-value of the null hypothesis given the data is the smallest significance level p for which any of the tests would have rejected the null hypothesis.

*Regression:* Linear regression fits a line to a scatterplot in such a way as to minimize the sum of the squares of the residuals. The resulting regression line, together with the standard deviations of the two variables or their correlation coefficient, can be a reasonable summary of a scatterplot if the scatterplot is roughly football-shaped. In other cases, it is a poor summary. If we are regressing the variable Y on the variable X, and if Y is plotted on the vertical axis and X is plotted on the horizontal axis, the regression line passes through the point of averages, and has a slope equal to the correlation coefficient times the SD of Y divided by the SD of X.

*Standard deviation (SD)*: The standard deviation of a set of numbers is the root mean square of the set of deviations between each element of the set and the mean of the set.

*t-statistic*: A statistic used to test hypotheses. A hypothesis test can be constructed by deciding to reject the null hypothesis when the value of the test statistic is in some range or collection of ranges. To get a test with a specified significance level, the chance when the null hypothesis is true that the test statistic falls in the range where the hypothesis would be rejected must be at most the specified significance level.

## Endnotes

[1] Daniel Lazier is a search fund entrepreneur and the founder of Skyline Crest Partners, a California-based search fund.

[2] Jacob Thomas is the Williams Brothers Professor of Accounting and Finance at the Yale School of Management.

[3] A. J. Wasserstein is the Eugene F. Williams, Jr., Senior Lecturer in the Practice of Management at the Yale School of Management.

[4] Kelly, Peter, and Heston, Sara. *2024 Search Fund Study Selected Observations.* Stanford University Graduate School of Business. Case E-870.