Distribution Truth and Consequences in the Forecasting of Headcount Growth

1. Introduction

Forecasting the growth rate of an organization's headcount is sufficiently difficult that real estate strategy is frequently constructed on the basis of flat growth or simple compounded growth models. By ignoring this issue or the arguments that might stem from constructing a more sophisticated model we are assured of only one thing: our forecasts will be uninformed and uninformative.

The primary reason cited for using flat growth (for example) is that external events (rather than more predictable internal plans) drive growth and these external events (mergers or acquisitions or sudden changes in the fortune of the core business) cannot be predicted. While there is an element of truth in this observation, it does not justify abdication; merely a more sophisticated approach.

This paper outlines such an approach, borrowing ideas from Wall Street and thinkers like Benoit Mandelbrot (Mandelbrot & Hudson, 2004) and Eugene Fama (Fama, 1965), whose independent analyses of financial market data both pointed toward the inadequacy of the normal distribution. It also follows closely the concepts underlying the value at risk (VAR) approach practiced by portfolio managers and actuaries (Jorion, 2007). It attempts to shed light on the occurrence of unlikely events ("black swans" (Taleb, 2008)) and the financial consequences of those events. The modeling of these consequences can help to answer whether these more sophisticated approaches do a better job and if so, how much of a better job.

2. Plan for the Paper

After a review of the applicable literature, the paper will describe a dataset of employee headcount of 20 technology firms. This data will be processed to create statistics on the rates of growth of these firms. This growth rate meta-data will be the main subject of the analysis. The core of the analysis will be to fit those rates of growth to two distributions (normal and stable), and to test how well those distributions fit the data using several standard goodness-of-fit statistics, as well as several visualizations.

While these tests will show that the stable distribution family fits the data "better," it will not show precisely how much better, and will therefore not provide guidance as to whether using the stable distribution family produces a materially better analytical outcome.

The "Swiss Army Knife" for these types of problem is the Monte Carlo simulation, which allows us to use an input distribution to condition the evolution of the process, and which returns an output distribution that can be examined at any percentile of interest, for example the 75th or 99th percentile outcomes. In order to determine the materiality of the improvement, two different Monte Carlo simulations will be run. These simulations will model future headcount growth using each of the distributions, and will run the output through a "consequence model" that will quantify the results of choosing each of the distributions.

The consequence model provides a simplified view of the costs of creating a real estate growth plan (focusing on leased space) that attempts to meet the needs of future growth. The model quantifies the consequences of taking too much or too little space under the conditions of simulated future growth. Because there is limited academic literature of the consequences of this misspecification, simplifying modeling assumptions are employed to construct a representative outcome.

This plan is illustrated in Figure 1:

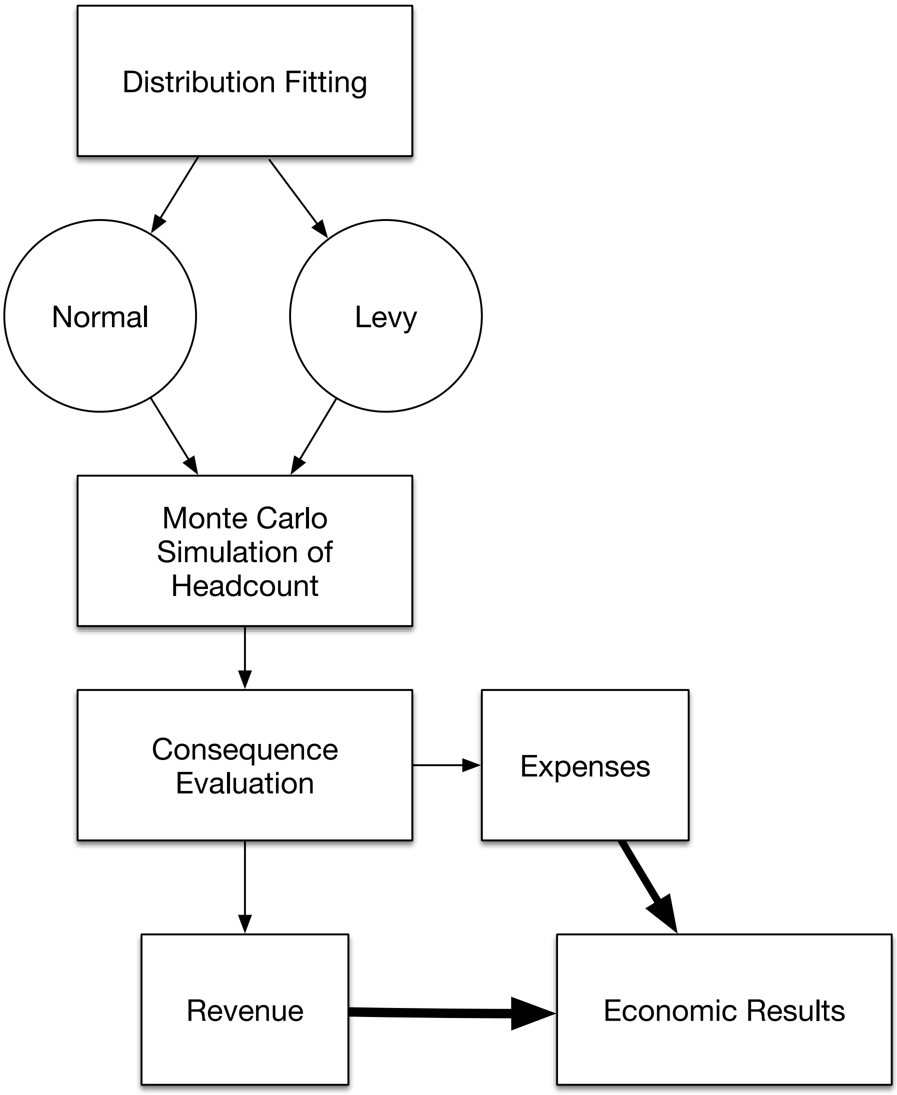

Figure 1

3. Background and Literature

Stable Distributions

Stable distributions are a class of four-parameter distributions that allow for skewness and for heavy tails. Benoit Mandelbrot proposed the use of these distributions in the context of financial markets. Mandelbrot, after attempting to fit changes in stock price data to the normal distribution and finding the fit poor, stated:

> *"..these facts warrant a radically new approach to the problem of price variation."(Mandelbrot & Hudson, 2004) p. 395*

Fama, working independently at the same time, reached a similar conclusion through empirical analysis of the Dow Jones index, finding that stock price changes exhibited far more extreme movements than the normal distribution would predict — a result he attributed to the fat-tailed character of the underlying distribution (Fama, 1965).

John Nolan, writing in his book *Univariate Stable Distributions: Models for Heavy Tailed Data* reviews the conclusions of prior researchers, stating that:

> *"It is argued that some observed quantities are the sum of many small terms–the price of a stock, the noise in a communication system, etc. and hence a stable model should be used to describe such systems." (Nolan, 2020) p. 24*

Nolan goes on to make the empirical observation that "many large data sets exhibit heavy tails and skewness" and this is true of the data employed in this analysis, as will be described in greater detail in the following section. Nolan addresses the common objection that stable distributions are impractical because they lack finite integer moments, arguing that the quality of the fit to the data is more important than this theoretical limitation (Nolan, 2020, p. 50). One of the practical challenges with stable distributions is a lack of standardization of the parameterization. "There are multiple parameterizations for stable laws and much confusion has been caused by these different parameterizations" (Nolan, 2020, p. 5). This paper employs Mathematica, and its stable-distribution generation functions[1].

Productivity and Office Space

One of the central issues in this analysis involves the relationship between productivity and office space. The basic premise of this paper is that there is an economic consequence associated with mis-estimating the amount of space a firm requires. The consequence of taking more space than needed is fairly straightforward: the firm pays more rent than it would have had to if the forecast was accurate.

The consequence of under-predicting demand is far more complex, and the analysis would be made much more concrete if there were empirical studies of lost revenue because of space, quantitative measures of lost productivity, or data on the impact of mis-predicting on levels of interior construction ("churn").

Unfortunately the literature does not provide clear direction on the financial consequences of overcrowded space. As Kaczmarczyk and Murtough noted in their 2002 article *Measuring the performance of innovative workplaces*:

> *"..the problem of productivity measurement of knowledge (formerly office) work has been puzzling researchers and practi-*

1. Specifically Mathematica's StableDistribution[] function and the S(0,α,β,γ,δ) parameterization. There is a difference in order between the Mathematica third and fourth parameters and those described in Nolan; the function of the parameters (location and scale) is the same.

tioners for at least 20 years."(Kaczmarczyk & Murtough, 2002) p. 168

Reviewing the efforts of the GSA to quantify white collar productivity, the authors conclude:

> *"The GSA research team realised that attempting to solve the problem of how to measure the productivity of knowledge workers was in itself a task that was unproductive (so to speak) and probably impossible in its purest form."(Kaczmarczyk & Murtough, 2002) p. 178*

The majority of studies have focused on employee satisfaction with their environment (through a self-reported survey-based methodology) or through analysis of the effect of HVAC and indoor air quality, on productivity. See for example, Roelofsen.(Roelofsen, 2002)

The conclusions of many of these studies are represented by this quote from Leaman describing the results of a study of 5,000 workers conducted by the British consultancy Building Use Studies. The survey asked: "Please rate how you think the physical conditions at work influence your productivity." Leaman concludes:

> *"as it is impossible to measure productivity satisfactorily across all types of managerial, professional and clerical work in offices of different kinds, this type of question serves as a useful, if methodologically weak, indicator."(Leaman, 1990) p. 12*

The study did show, however, that 53.1% of office workers surveyed felt that a bad workspace reduced their productivity by 10% or more.*(Leaman, 1990)* p. 13.

Haynes, in *The impact of office comfort on productivity* quotes Mawson 2002 (Mawson, 2002):

> *The challenge to find a validated method of measuring and reporting office productivity remains to be achieved, with some authors referring to this area of research as the "search for the Holy Grail"(Haynes, 2008) p 39.*

While these types of studies provide useful information about the relationship between design and work, they do not provide specific data for:

- the effect of overcrowding on white collar output or revenue;
- the effect of space shortages on levels and costs of churn; or
- the business opportunities (and revenue) forgone by growing firms resulting from a dearth of space.

4. Description of Data

The data analyzed for the paper was derived through a combination of SEC filings and Moody's data. It consists of headcount data series for 20 technology firms. The firms include software, web, new media, and hardware.

In order to get a general assessment of the rates of growth of these firms, the following data manipulation was performed. First, 660 data points were log transformed. Second, the differences between logs was computed. Then that difference was exponentially transformed. Increases of greater than 2x between any two years were filtered, on the assumption that these outliers represented early period growth at very small headcount numbers or incorrect data. This resulted in 625 data points.

The right-skewed, fat-tailed character of this data is consistent with the broader empirical literature on firm growth. Michael Stanley and his associates analyzed publicly traded US manufacturing companies over a seventeen-year period, and found that the distribution of annual growth rates departs significantly from the normal, exhibiting an exponential form with heavier tails than standard models would predict, a finding that has since been replicated across multiple industries and geographies (Stanley et al., 1996). This is illustrated in the following histogram:

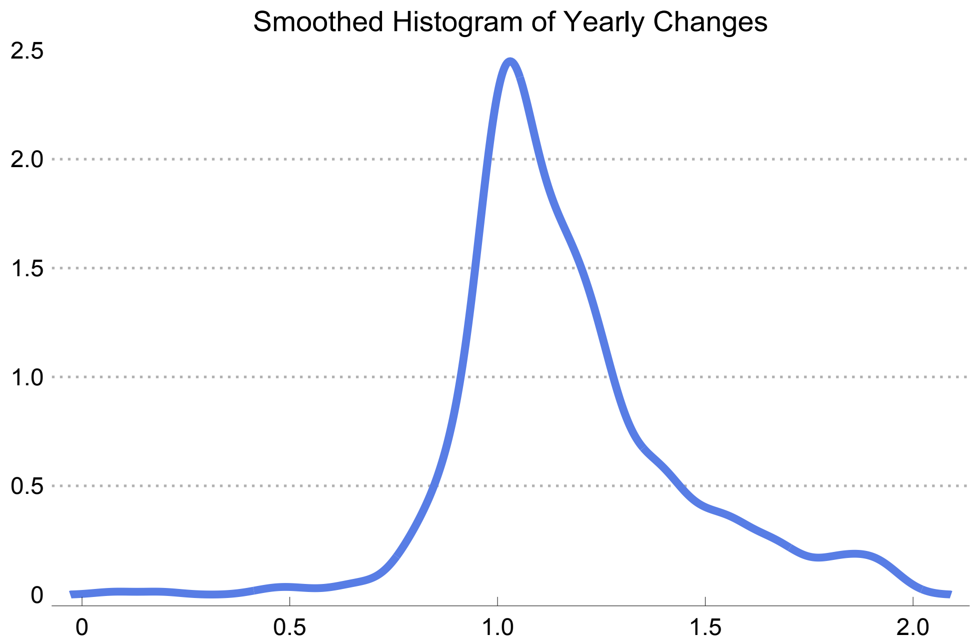

Figure 2a

To understand the average growth rates of each company, a second data manipulation was performed, by averaging the differences in the logs. These averages were then exponentially transformed to give an average yearly growth rate for each firm. Figure 2b illustrates these growth rates for each of the firms. Values range between .912 (WebMD, which experienced dramatic growth followed by an equally dramatic downsizing) through 1.541 (Yahoo). The average of the averages was 1.24.

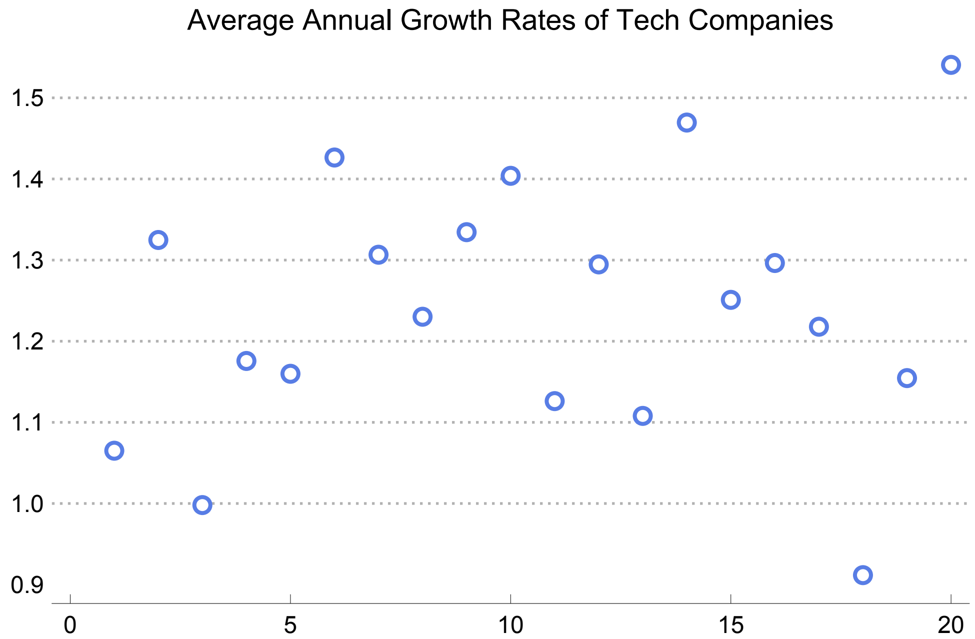

Figure 2b

5. Fitting Distributions to the Data

The dataset was fitted to a variety of distributions, including the normal, the log-normal, and Cauchy, and to the "0" parameterization of the stable distribution. The following Figure 3 shows the histogram of the data compared with the best-fit parameterization of these four distributions.

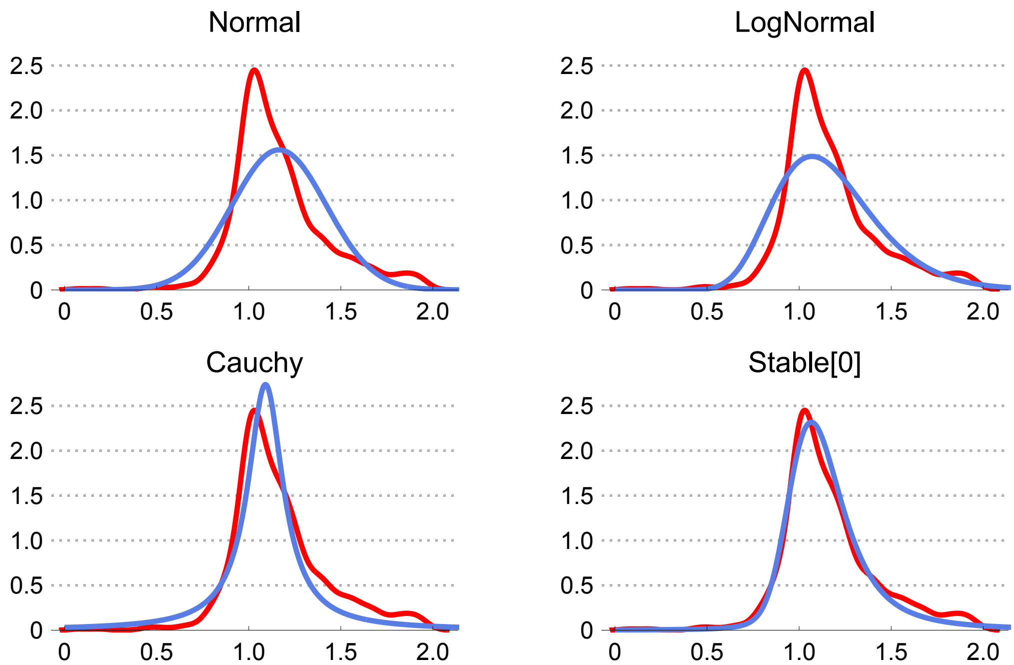

Figure 3

Since the data is clearly skewed to the right and fat-tailed, log-normal, Cauchy and stable distributions have the best visual fit. To focus in on the central question proposed at the beginning of the paper, Figure 4a provides two probability plots of the data, one illustrating the fit to the normal and the other illustrating the fit to the stable distribution. It is evident from this comparison that the stable distribution is a better fit.

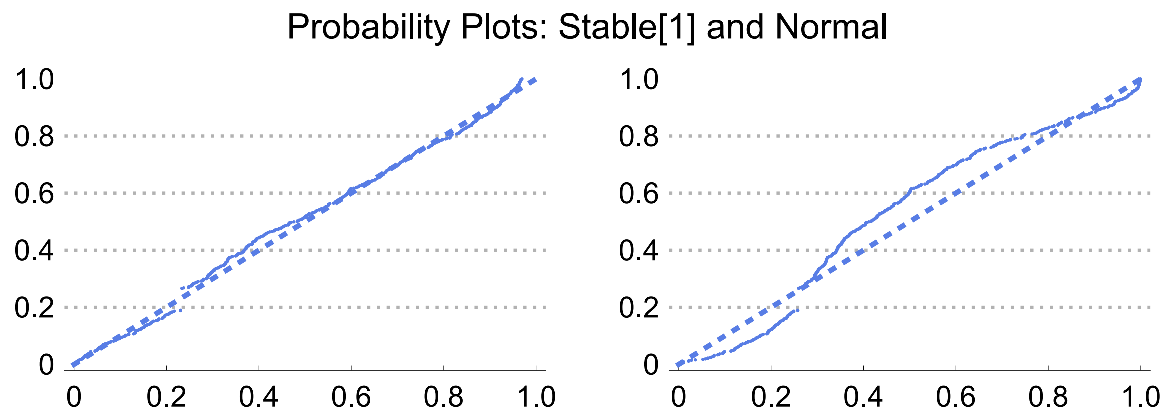

Figure 4a

While the PDF to histogram comparison makes the Cauchy distribution look like a good candidate, a q-q plot shows significant deviation at the extremes of the distribution, as can be seen in Figure 4b.

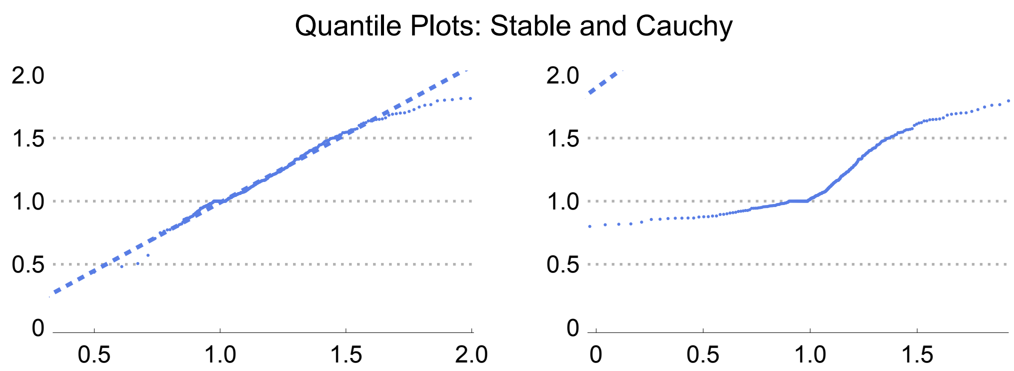

Figure 4b

Unfortunately most statistical tests of fit are rendered invalid once the parameters of a distribution have been estimated from the data.

6. Consequence Model Description

The correct approach for constructing a consequence model would be to make a "bespoke" model for each tenant and market situation. The consequences would include both lost revenue and increased cost, as in Figure 5:

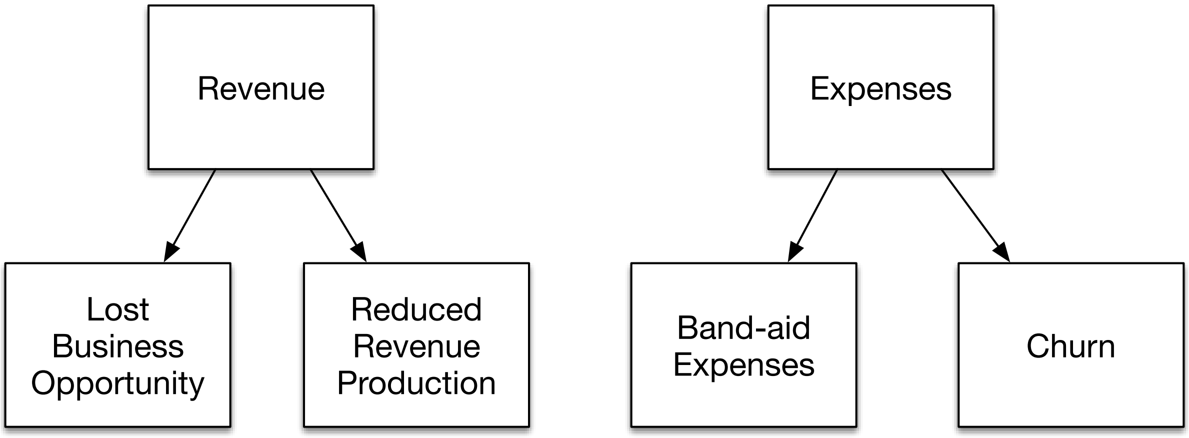

Figure 5

This model would have to consider:

1. Revenue consequences of (potentially) *foregone business opportunities*. In other words lines of business that would have been created but were stunted by the lack of space. There does not seem to be any data on this subject in the literature; this would have to be modeled as a rare but highly consequential event.

2. Revenue consequences of *lost productivity*. This category would have two sub-categories. The first would be productivity declines from over-crowding: in other words less than optimal performance by personnel in a revenue-creating category (for example salespeople) resulting from inappropriate space. The second would be "flight risks": employees in a revenue-creating capacity that leave the organization in part or in large part because of inadequate facilities.

3. *Band-aid costs.* These would include the increased cost of space and fit-out associated with un-planned space acquisitions. These costs would result both from the need to build space rapidly and a rent premium on the cost of leasing space outside of a well-structured plan.

The actual cost of band-aid space will be highly market-specific. It will frequently be more expensive to take band-aid space than the "main" space, since it would likely be for sub-optimal term (matching the main space), and since as a compromise it will be often necessary to "pay up" to minimize disruption to broader business operations. In some instances the creation of "A" and "B" space (i.e. primary space where the core of business operations go on and secondary, less desirable space where non-core activities get consigned) will require additional costs to make the "B" space more palatable to groups compelled to operate in it.

4. *Churn and swing space*. There are methods to cope with operating with less than optimal amounts of space (particularly at the margin and before band-aid space becomes critically necessary). These methods include doubling-up, recycling of common space (for example conference rooms), and densification strategies (benching e.g.). Most of these techniques have in common the need to spend capital to construct or modify space. To perform these modifications, the firm must have a strategy for staged relocation, which often takes the form of swing space areas to house groups while construction proceeds. An organization at the margin of over-capacity could need 10% (or more) of its space to be designated as swing space.

For the purpose of this paper we will not attempt to construct such a bespoke model of consequences. We will instead rely on a simplification which will be a quadratic model in which space needs incremental to that initially contracted will be modeled to cost the square of the increment. Therefore additional space amounting to 20% beyond the initial envelope (1.2 times) will be modelled to cost 44% more (1.44 times) the cost of the initial space. This quadratic formulation will cause large deviations from the plan to have disproportionate impacts on the cost. While this simplified model will not capture many of the impacts described above, it will be useful to assess whether there is a meaningful difference between growth modeled as a sequence of normally distributed changes, versus growth modeled as a sequence of stable-distributed changes.

7. Monte Carlo Simulations

The core of the Monte Carlo simulation is very straightforward. A growth rate is selected randomly from the chosen parameterized distribution. This growth rate is applied, and this process is repeated for a number of periods. The result is then input to the consequence model, which provides a non-linear (quadratic) transformation to the headcount data to convert it into a cost of under or oversupply of space. This process is repeated and the results collected and analyzed as a distribution. Figure 6 illustrates the simulation:

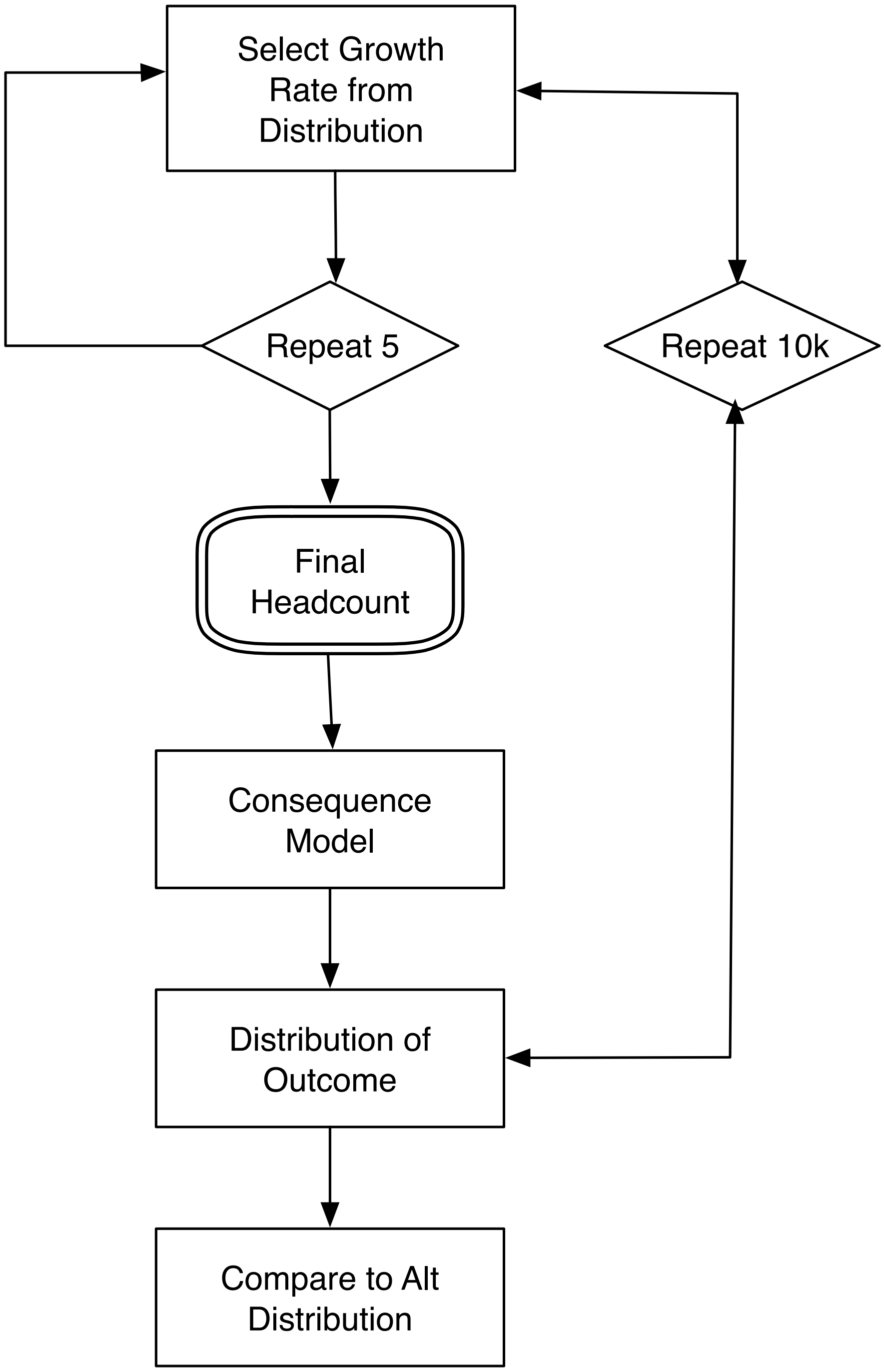

Figure 6

8. Summary of Simulation Results

The simulation described in Section VII was run for ten thousand trials over a five-year forecast horizon, a period consistent with a standard technology company lease term. The costs reported below represent the final-year consequence cost for each trial--that is, the cost associated with the gap between the space contracted at the outset and the space actually required at the end of the planning horizon.

Before examining costs, it is useful to consider the growth rate distributions that the simulation produces. After five compounded steps, the normal distribution yielded a median cumulative growth of 1.95 times the initial headcount; the stable distribution yielded 1.97 times: for practical purposes indistinguishable. The means diverge more: 2.15 times initial headcount for the normal versus 2.71 times for the stable. But the critical difference is in the standard deviations: 1.10 times initial headcount for the normal and 5.08 times for the stable, a ratio of more than four and a half to one. At the 99th percentile, the stable distribution produced cumulative growth of 14.1 times the initial headcount, compared to 5.6 times under the normal. The maximum simulated growth across all 10,000 trials reached 263 times the initial headcount under the stable distribution, versus 9.5 times under the normal.

These differences in the tails of the growth rate distribution matter because of the quadratic cost structure described in Section VI. Modest forecast errors produce modest cost penalties; large ones produce disproportionately large ones. The central question is whether the heavier tail of the stable distribution,which has only a limited effect at the median, produces materially different costs when the full distribution of outcomes is considered.

Normal Distribution

Under the normal distribution, the five-year simulation produced a median final-year consequence cost of $5.6 million. The mean was $6.5 million, reflecting the right skew introduced by the quadratic cost function. The standard deviation was $4.0 million. At the 90th percentile the cost reached $11.4 million; at the 95th percentile, $13.8 million; and at the 99th percentile, $21.4 million — a ratio of 3.8 to one relative to the median. The maximum observed cost across all trials was $42.4 million.

Looking at how extreme costs compound over time, the 99th percentile cost under the normal distribution grew from $5.0 million at the end of year one to $7.4 million at year two, $10.0 million at year three, $14.3 million at year four, and $21.4 million at year five. The compounding is orderly (approximately 1.4 times per year at this percentile) and broadly consistent with the growth rate assumptions underlying the simulation.

Stable Distribution

The stable distribution produced a median final-year consequence cost of $5.6 million, which is identical to the normal distribution. This result is not coincidental. It reflects the fact that the consequence model operates in discrete 5,000 square foot increments and that the central mass of both distributions passes through the same set of lease-decision thresholds in a similar way. At typical growth rates, the two distributions produce the same real estate decisions and therefore the same costs. A planner using either distribution as the basis for a five-year space plan would arrive at the same answer for the most likely case.

The divergence begins in the upper percentiles and accelerates sharply. At the 75th percentile the stable distribution produces a cost of $8.5 million versus $8.1 million under the normal — a ratio of 1.05, barely distinguishable. At the 90th percentile the gap widens: $14.3 million versus $11.4 million, a ratio of 1.25. At the 95th percentile the difference becomes practically significant: $23.8 million versus $13.8 million, a ratio of 1.72. And at the 99th percentile the stable distribution produces a consequence cost of $67.0 million — 3.14 times the $21.4 million produced by the normal distribution. The maximum observed cost across all 10,000 trials was $1.41 billion, against $42.4 million under the normal.

The year-by-year progression of the 99th percentile cost illustrates how quickly the distributions diverge as the horizon lengthens, as seen here in Figure 7:

Year	Normal 99th Pct	Stable 99th Pct	Ratio
1	$4.95M	$8.88M	1.79
2	$7.38M	$17.75M	2.41
3	$10.00M	$32.15M	3.22
4	$14.28M	$47.75M	3.35
5	$21.35M	$66.95M	3.14

Figure 7

By year three the ratio has stabilized in the range of 3.1 to 3.3 and remains there through year five. This is an important observation: the tail risk premium associated with using the correct distribution is not open-ended or chaotic. It is a persistent and compounding multiplier on worst-case outcomes.

The mean cost under the stable distribution was $9.4 million, 43% higher than the normal distribution's $6.5 million mean. Unlike the median, the mean is sensitive to the heavy tail and therefore serves as a rough indicator that the expected cost of being wrong about the distribution is material even when averaged across all outcomes, good and bad.

9. Conclusion

The analysis presented in this paper set out to answer a practical question: does it matter, in a commercial real estate context, whether the distribution used to model headcount growth is a normal distribution or a member of the stable family? The answer, based on the data examined and the simulation constructed, is that it matters enormously, but only some of the time, and only in one direction.

The data from 20 technology companies, covering 625 annual growth observations, does not fit the normal distribution well. As Figure 3 illustrates, the data is right-skewed and fat-tailed, and the stable distribution provides a materially better visual fit at the tails that matter most for planning purposes. Figures 4a and 4b confirm this: the probability plot for the stable distribution tracks the data closely, while the normal distribution systematically understates the frequency of both growth and contraction at the extremes.

The consequence simulation demonstrates that this distributional mis-specification is not merely a statistical curiosity. A planner who uses the normal distribution will, in the typical case, construct the right real estate strategy. The median outcome is $5.6 million under both distributions. The 75th percentile outcomes are within 5% of each other. For most tenants, in most years, the difference between the two approaches is negligible. This is precisely why the problem is easy to overlook: the normal distribution performs adequately in the center of the distribution, where outcomes are measured most often and where intuition is calibrated.

The problem arises at the margins. At the 90th percentile the stable distribution produces predicted costs 25% higher than the normal. At the 95th percentile the premium reaches 72%. At the 99th percentile the stable distribution produces predicted costs more than three times those predicted by the normal. These are not tail events in any exotic sense; with a portfolio of twenty properties, there is roughly a 64% chance that at least one will exceed the 95th percentile cost in any given year. For a tenant managing multiple locations, or for a landlord pricing renewal risk across a large portfolio, the difference between a 95th percentile cost of $13.8 million and $23.8 million is a material planning error.

The year-by-year analysis adds a further dimension. The divergence between the distributions is largest and most consequential at longer time horizons. After one year the 99th percentile ratio is 1.79; by year three it has risen to 3.2, where it remains through year five. A standard tech company lease term of, say, five years therefore coincides with exactly the horizon over which the normal distribution most severely underestimates tail risk.

While the data examined here comes from technology firms, there is no reason to believe the finding is sector-specific. Any commercial tenant whose headcount growth is subject to external shocks, whether from market cycles, acquisitions, product pivots, or macro-economic disruption, is likely to exhibit a growth distribution with heavier tails than the normal. The techniques employed here, fitting growth rate data to the stable distribution and running consequence simulations, are directly transferable to other tenant categories: financial services firms, law firms, healthcare organizations, or any occupier whose space demand is driven by a headcount that is itself driven by uncertain future events.

As we move forward and our analytical tools become more sophisticated and more accessible we should be using tools that characterize uncertainty that match the nature of that uncertainty. The value at risk framework borrowed from portfolio management--fitting distributions, running simulations, examining percentile outcomes rather than point estimates--provides a structured and reproducible method for doing so. The result is not a single number but a distribution of outcomes, with the tails properly weighted. That is, to borrow Mandelbrot's formulation, a less misbehaved approach to forecasting headcount growth, a problem that has been misbehaving for some time.

Bibliography

Leaman, A. (1990). Productivity and office quality. Facilities, 8(4), 12-14. Haynes, B. P. (2008). The impact of office comfort on productivity. Journal of Facilities Management, 6(1), 37-51. doi:10.1108/14725960810847459

Fama, E. F. (1965). The behavior of stock-market prices. The journal of Business, 38(1), 34-105.

Jorion, P. (2007). Value at Risk: The New Benchmark for Managing Financial Risk. McGraw-Hill.

Mandelbrot, B. B., & Hudson, R. L. (2004). Misbehavior of Markets. Basic Books.

Mawson, A. (2002). The workplace and its impact on productivity. Advanced Workplace Associates, London, 1, 1-12.

Nolan, J. P. (2020). Univariate stable distributions. Springer.

Roelofsen, P. (2002). The impact of office environments on employee performance: The design of the workplace as a strategy for productivity enhancement. Journal of Facilities Management, 1(3), 247-264. Retrieved from http://dx.doi.org/10.1108/14725960310807944

Kaczmarczyk, S., & Murtough, J. (2002). Measuring the performance of innovative workplaces. Journal of Facilities Management, 1(2), 163-176. doi:10.1108/14725960310807908

Stanley, M. H. R., Amaral, L. A. N., Buldyrev, S. V., Havlin, S., Leschhorn, H., Maass, P., . . . Stanley, H. E. (1996). Scaling behaviour in the growth of companies. Nature, 379(6568), 804-806.

Taleb, N. N. (2008). The Black Swan: The Impact of the Highly Improbable. Penguin Books Limited.