



Economies of extremes: Lessons from venture-capital decision making



Suzanne de Treville^{a,*}, Jeffrey S. Petty^a, Stefan Wager^b

^a University of Lausanne, Faculty of Business and Economics, 1015 Lausanne, Switzerland

^b Stanford University, Department of Statistics, Stanford, CA 94305, USA

ARTICLE INFO

Article history:

Received 5 September 2013

Received in revised form 24 June 2014

Accepted 10 July 2014

Available online 27 July 2014

Keywords:

Capacity strategy

Extreme-value theory

Venture-capital decision making

ABSTRACT

An organization's ability to exploit extreme events—such as exceptional opportunities—depends on its capacity strategy. The venture capital industry illustrates the interplay of expensive capacity and negative externalities from high utilization. The cost of adding a venture capitalist provides a strong incentive to run lean, but such leanness may make it impossible to evaluate all interesting investment opportunities. Using concepts from extreme-value theory, we analyze the trade-off between the costs and benefits arising from an increase in the number of evaluated deals. We ground our analysis in 11 years of archival data from a venture capital firm, representing 3631 deals, the decisions made, the reasons for those decisions, and the decision lead times. The firm identified 20% of arriving deals as worth evaluating during the screening process, but was not able to evaluate approximately 9% of those interesting deals due to a lack of capacity. We show that the value of increasing the number of deals evaluated increases with the tail weight of the distribution of deal values. When the right tail is light, increasing the number of deals evaluated may provide too modest a benefit to justify the cost. When, however, the right tail is heavy, the value of increasing the number of deals is likely to more than compensate for the cost of capacity. Our results provide new insight into the relative value of a chase capacity strategy that emphasizes responsiveness versus a high-utilization heuristic that emphasizes productivity. Our approach can be applied to other search operations such as personnel selection, quality circles seeking to identify root causes, and making employee capacity available for innovation.

© 2014 Elsevier B.V. All rights reserved.

1. Introduction

In a survey of literature on strategic capacity management, Van Mieghem (2003) described the tension arising from the interplay of expensive capacity and negative externalities from high utilization, noting that unmet demand is likely when capacity is expensive. In this paper, we use a search operation to study how extreme values shape this interplay. Our analysis is grounded in 11 years of historical data concerning capacity shortages occurring in a venture-capital firm. During the two funds covered by the data, the firm followed a policy of only adding capacity when absolutely unavoidable because of the high cost of hiring an additional venture capitalist. When a deal arrived to the venture-capital firm, an initial screening selected those deals deemed interesting enough for a thorough evaluation. Of deals recorded by the decision maker as interesting enough for evaluation, 9% were

nevertheless rejected, with the reason cited being lack of capacity. The historical data thus suggests a type of level capacity strategy in which demand was foregone when the maximum service rate was exceeded. The research question that arose from the historical data was whether the savings from postponing an additional hire—running lean—compensated for the cost of not evaluating all interesting deals.

The increase in fund value arising from an increase in capacity allocated to evaluation depends on many factors, including the quality of the additional deals evaluated, the competence of the venture capitalists in choosing deals, the dynamics of the deal-arrival process, and the heaviness of the right tail of the distribution of deal values. Our focus is the right-tail weight. A venture-capital firm invests in a small number of deals relative to the sample from which they are selected: The firm described here received 3631 deals, evaluated 700, and selected what they considered to be the best 35 of those 700. Had they evaluated more than 700 deals, some deals at the bottom of the top 35 out of 700 probably would have been replaced by better deals. Does swapping out a small number of good deals for better deals increase the overall

* Corresponding author. Tel.: +41 21 692 3448.

E-mail address: suzanne.detreville@unil.ch (S. de Treville).

portfolio value enough to justify the cost of the extra capacity required?

Extreme-value-theory methods emerge as natural tools for addressing this question. As we deepen the search, the value of a portfolio made up of the top 35 deals will increase at the same rate as the sample maxima. We can characterize this rate under regularity conditions. We show that the relationship between capacity strategy and profit hinges on the magnitude of extremes, as captured by the tail index of the distribution of deal values. The tail index is an extreme-value-theoretic construct that describes the decay rate of the right tail. Distributions with a thin right tail, such as the Gaussian distribution, have a null index. As the tail index approaches 0.5, the variance becomes infinite, and at a tail index of 1 the mean becomes infinite.

To make the problem mathematically tractable, we begin with three simplifying assumptions: (1) the venture capitalists are perfectly competent at evaluating deals, (2) all deals are considered simultaneously, and (3) the additional deals evaluated are drawn from the same distribution as the original deals. We show that for a tail index of γ , an increase in the number of evaluated deals by a factor of α is expected to increase the average value of the portfolio by α^γ . We later demonstrate that this result is robust to violations of the first two simplifying assumptions, and that it is possible to estimate the impact of violating the third assumption.

To gain intuition about the tail weight of deal values that a venture-capital firm might be expected to face, we estimated the tail index of the distribution of returns of initial public offerings (IPOs). If the tail weight of deal values resembles the $\gamma=0.4$ tail weight of IPO returns, our model predicts that a 9% increase in the number of deals evaluated should produce an expected increase of 3.5% in the value of a portfolio made up of the 35 top deals. Applied to an average venture-capital portfolio of \$151 million, a 3.5% increase yields \$5 million, more than enough to fund an increase in capacity.

Our results extend beyond the venture-capital context, and contribute to a better understanding of the value of capacity in a search operation. The general insight that increasing sampling by a factor of α is expected to increase the value of the sample maxima by a factor of α^γ can be used to quantify the value of allocating capacity to a search operation. The tail weight has been incorporated as one parameter of the distribution of concept values in balancing cost and benefit of search in product design (Dahan and Mendelson, 2001). Tail weight can also play a role in the decision about how much capacity to allocate to search operations such as personnel selection and root-cause identification by quality circles. In Section 6 we describe the application of the model to a decision faced by a large organization about how much capacity to allocate to personnel selection.

More generally, our analysis gives insight into how a positive tail weight impacts the value of a capacity buffer. Having a capacity buffer—that is, following a chase capacity strategy—allows a firm to take advantage of peak demand. Managers often find it difficult to justify maintaining slack capacity just in case peak demand would occur, especially in light of the lean-production paradigm and its focus on eliminating rather than accommodating variability (e.g., Womack et al., 1990). Firms seeking to maintain a high capacity utilization are more likely to follow a level capacity strategy that allows for stockouts and may entail building inventory during times of low demand. Deploying a level capacity strategy is often done in conjunction with practices aimed to tame demand variability. As extreme-value theory is not commonly used in the operations-management field, managers are typically not equipped to quantify the value of the capacity buffer, and may find it difficult to justify maintaining a capacity buffer even when intuition argues that it would be of high value. Van Mieghem (2003, p. 281) referred to the “continental divide between inventory and queuing models [that

also applies to the subfield of capacity strategy”. Incorporation of the tail index into the capacity-strategy literature may contribute to bridging this divide.

2. Relevant literature

One approach to capacity strategy arises from literature on reducing lead times, with its emphasis on the importance of avoiding excessive resource utilizations (e.g., Hopp and Spearman, 2001; Suri, 1998). The queuing models common to this literature provide a way of valuing a capacity buffer in the face of demand uncertainty. More generally, slack capacity has been established as a key source of volume flexibility (Jack and Raturi, 2002). Hayes and Wheelwright (1984) proposed as a rule of thumb that firms establish a modest capacity buffer (on the order of 10%) in the face of demand variability. Authors such as Kellogg and Nie (1995) and Goodale et al. (2003) have highlighted the need for capacity buffers in service operations facing high arrival-rate variability. Koste et al. (2004), however, observed that volume flexibility was less addressed in the literature than other types of flexibility, possibly indicating some ambivalence in how capacity buffers are viewed in the operations management field.

Van Mieghem (2003, p. 275) explored and summarized what he referred to as the “direct functional dependence of operating profit on the capacity stock” that depends on time, the resources that make up its “capacity portfolio”, and uncertainty. Uncertainty encourages firms to optimally plan demand-capacity imbalances. Van Mieghem (2003) observed; however, that those working to implement such a capacity buffer faced the hurdle of explaining to management the need to invest in capacity that might not be fully utilized.

The emphasis on maintaining high resource utilization is exemplified by the Toyota Production System and the lean-production approach that it spawned. Lean production emphasizes line balancing through practices such as bottleneck elimination and cycle-time reduction (e.g., Shah and Ward, 2003). Tools like heijunka combine with line balancing to allow utilization to be maximized (e.g., Womack et al., 1990). The capacity-profit functional relationship under lean production incorporates the idea that high utilization encourages learning, with intentional penalty of capacity used to create exploratory stress that encourages productivity improvements (e.g., Monden, 2011; Suri and de Treville, 1986).

Capacity strategy tends to focus on demand uncertainty. Extending the focus to search operations provides new insight into how to value a capacity buffer. Balancing the cost and value of search is addressed in literature about the innovation process (e.g., Chao et al., 2009; Dahan and Mendelson, 2001; Loch and Kavadias, 2002; Terwiesch and Loch, 2004; Terwiesch and Xu, 2008; Weitzman, 1979), with selection of deals by venture capitalists given as an example of such a search operation (Terwiesch and Xu, 2008; Terwiesch and Ulrich, 2009).

Dahan and Mendelson (2001) estimated the optimal number of concepts to test when the distribution of concept values follows a generalized-extreme-value distribution (GEV) with scale parameter b and tail-shape parameter $\alpha = \gamma^{-1}$, and with a per-test cost c . They showed that the optimal number of concepts to test for a null tail index was b/c , becoming $[(b/c\alpha)\Gamma((\alpha-1)/\alpha)]^{\alpha/(\alpha-1)}$ as the tail-shape parameter became positive. For a null tail weight, changes in the per-unit cost of testing were shown to have a constant effect on the optimal amount to spend on testing. As the tail weight became positive, the optimal spending for testing was shown to be convex increasing in decreases in c .

Although many of the researchers applying extreme-value theory to the innovation process have focused primarily on thin-tailed distributions (e.g., Erat and Kavadias, 2008; Girotra et al., 2010; Terwiesch and Xu, 2008), there is recognition that allowing more

متن کامل مقاله

دریافت فوری ←

ISIArticles

مرجع مقالات تخصصی ایران

- ✓ امکان دانلود نسخه تمام متن مقالات انگلیسی
- ✓ امکان دانلود نسخه ترجمه شده مقالات
- ✓ پذیرش سفارش ترجمه تخصصی
- ✓ امکان جستجو در آرشیو جامعی از صدها موضوع و هزاران مقاله
- ✓ امکان دانلود رایگان ۲ صفحه اول هر مقاله
- ✓ امکان پرداخت اینترنتی با کلیه کارت های عضو شتاب
- ✓ دانلود فوری مقاله پس از پرداخت آنلاین
- ✓ پشتیبانی کامل خرید با بهره مندی از سیستم هوشمند رهگیری سفارشات