Optimal Parallel and Sequential Prototyping in Product Design

Ely Dahan
Sloan School of Management
Massachusetts Institute of Technology
Cambridge, MA 02142-1347

Haim Mendelson
Graduate School of Business
Stanford University
Stanford, CA 94305-5015

October 1998

Optimal Parallel and Sequential Prototyping in Product Design

Abstract

When designing a new product, developers must balance the expected profit of improving the design against the cost and time requirements of doing so. We investigate this tradeoff by modeling the prototyping stage of product development as a probabilistic search process. We study optimal policies for single-period and infinite-horizon problems and derive closed-form solutions for the case of profits distributed according to the three extreme-value distributions. We demonstrate the role played by the tail-shape parameter of the distribution in determining the level and form of experimentation. The effect of declining prototype costs on expected profits, number of experiments, and spending is shown to depend on the form of the profit distribution and its tail-shape parameter. The paper compares the performance of pure parallel, pure sequential and hybrid parallel-sequential experimentation policies.

Optimal Parallel and Sequential Prototyping in Product Design

1. Introduction

Prototyping provides valuable information to designers about the customer utility, technical feasibility and cost of a new product. A customer-ready prototype enables designers to discover problems and reveal opportunities that might have been overlooked in the design's conceptual phase. One of the most challenging decisions facing design teams is allocating resources to tasks such as prototyping and testing, since their payoffs are highly uncertain.

We consider four modes of prototyping, which are summarized in Table 1:

- *One-shot*: The design team builds and tests *only one prototype*, consistent with a "Do it right the first time" philosophy.
- *Sequential*: The design team builds and tests *one prototype each period*, for as many periods as are necessary to achieve a "good enough" design.
- *Parallel*: The design team builds and tests *multiple prototypes in a single period*, then chooses the most profitable one.
- *Hybrid*: The design team builds and tests *multiple prototypes each period*, for as many periods as are necessary to achieve a "good enough" design.

Table 1: Prototyping Modes

Prototypes per period	Time-Frame: One period	Time-Frame: Multiperiod
1	One-shot	Pure Sequential
Many	Pure Parallel	Hybrid (Parallel and Sequential)

The one-shot prototyping approach conforms to the common wisdom that freezing design specifications early improves profits by coordinating the efforts of the functional areas of the firm, minimizing costly redesign iterations and speeding time-to-market. Given low technical and market risks, the one-shot mode maximizes design efficiency.

Figure 1

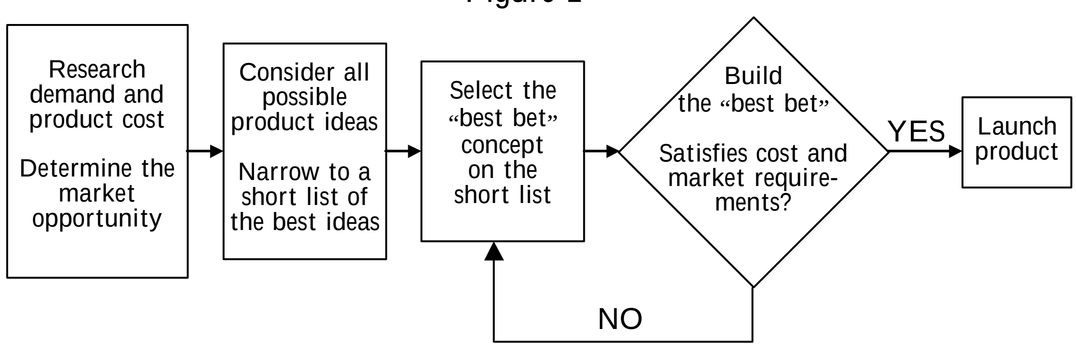

Figure 1 depicts the *sequential* mode of search. After researching the market and technology, a large number of product ideas are considered and narrowed down to a short list of leading concepts worthy of testing, as described by Ulrich and Eppinger (1995). The concepts on the short list are then ranked according to a process that includes competitive benchmarking, Pugh concept selection and similar structured methodologies. Next the "*best bet*" alternative is chosen as the first step in the search. This design alternative is prototyped and tested to see if it satisfies internal requirements for market demand, manufacturing cost and technical feasibility prior to being launched. An iterative process of improvement ensues until an observed outcome meets or exceeds the internally imposed profit hurdle.

For products with long design cycles, the sequential-iteration approach is effective. However, an optimal prototyping policy should consider not only development cost and product performance, which favor sequential prototyping, but also the benefits of faster time-to-market, which favor parallelism. Due to shrinking product life cycles and competitive pressure, the design team may need to accelerate the process of testing promising prototypes (c.f. Fine 1998, Mendelson and Pillai 1998). An empirical study by McKinsey claimed that the penalty for being six months late in releasing a new product is about a third of overall profits [Vesey 1991]. Products such as Internet software, computers, pharmaceuticals, semiconductors and even movies, to name only a few, demand urgency in their design and development. The expression "time is money" is literally true in these situations.

Accelerating clockspeed increases demand for parallel prototyping since product development is speeded up as shown in Figure 2. As before, the market and technology are researched, and then a plethora of ideas within the product category are considered. In the parallel mode, however, leading design concepts need not be narrowed to a single best concept. The design team decides how many parallel prototypes to build by balancing the cost of prototyping against the benefit of expected profit enhancement, given a single period in which to conduct experiments. After prototyping, the firm launches the design with the highest profit projection or abandons launch if even the best design is projected to lose money.

Figure 2

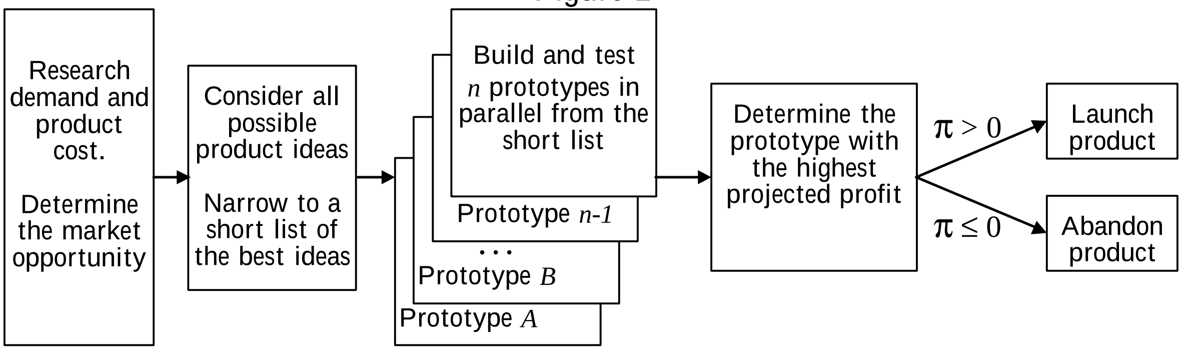

On the supply side, declining prototyping costs improve the economics of parallel prototyping. Virtual and rapid prototyping techniques portend dramatic declines in the cost per prototype. Virtual prototypes, such as those employed by Boeing for its 777 aircraft, allow design concepts to be tested without the expense and time requirements of physical prototyping. When distributed over the World Wide Web, using formats such as VRML (virtual reality markup language) or streaming video, virtual prototypes allow potential customers to experience the "look and feel" of new designs before physical versions are built [c.f. Dahan and Srinivasan 1998]. Preference and demand data may be collected efficiently in this way.

Combinatorial prototyping methods [Thomke, et. al. 1997] enable researchers to efficiently search through thousands of variations of custom-designed chemical compounds, to see, for example, which metallic oxide is the best superconductor or which organic compound is the most pharmacologically active. Rapid prototyping allows computer-controlled technology to create topologically complex physical prototypes using layering and etching techniques in place of traditional methods such as machining and casting. The economics of getting "from art to (prototyped) part" continue to improve, making parallelism more affordable.

In summary, the attractiveness of parallel prototyping is enhanced by the increased importance of time-to-market, the steady decline in prototyping costs, and the availability of new prototyping technologies.

Prior research related to sequential and parallel prototyping can be broadly divided into three categories: (1) models of search, (2) prototyping processes, and (3) R&D projects as real options.

Models from the broad literature on optimal search provide insight into the economics of prototyping. Weitzman [1979] develops an optimal solution to the problem of sequential experimentation under uncertainty, but does not consider parallel search. This is natural in the context of his analysis, since when *individuals* search for low prices or high-paying jobs, parallel search is not feasible. Weitzman's optimal policy ranks uncertain alternatives by their reservation prices, which depend on the reward distribution and cost of search. Weitzman orders the experiments and delineates the optimal stopping rule for a pure sequential policy (discussed in section 3.1).

Morgan and Manning [1985] study the optimization of sequential and parallel search policies and prove that a hybrid parallel/sequential policy dominates both pure policies. Gross [1972] extols the wisdom of developing multiple advertising campaigns and develops a one-period model of optimal parallel search incorporating residual uncertainty after the advertising campaigns are designed and tested. The number of parallel concepts to be tested is determined heuristically.

Srinivasan, Lovejoy, and Beach [1997] provide empirical evidence that parallel prototyping resolves some of the residual uncertainty remaining after the concept phase of product design. They show that, under reasonable assumptions, parallel prototyping is more profitable than the one-shot approach and propose parallel development of customer-ready prototypes as an attractive method of reducing risk.

A second line of research suggests that parallel prototyping requires appropriate processes, resources, and organizational structure. In a rapidly changing market environment, the design process must facilitate high-speed, high-performance innovation. Smith and Reinertson [1995] emphasize the importance of speeding up new product development and suggest parallelism as one method for doing so. Leonard-Barton [1995] coins the term "failing forward" to describe experimentation in which firms learn from failures and take advantage of the creativity inherent in highly variable experimental outcomes. Wheelwright and Clark [1992] discuss the positive and negative organizational implications of having internal teams compete to design the same new product. Thomke, Von Hippel, and Franke [1997] demonstrate how new prototyping technology facilitates parallel search. They cite, as an example, combinatorial chemistry in which pharmaceutical firms conduct large numbers of parallel experiments in the search for new drugs.

The third line of research suggests that multiple prototypes be viewed as *real options*. Hauser [1996] provides a comprehensive annotated bibliography on valuing R&D projects, including research on their real options nature. For example, Hauser and Zettelmeyer [1997] characterize optimal policies for managing R&D portfolios and show the value of the options provided by alternative solutions to a given problem, even if the alternatives come from R&D spillovers from outside firms.

This paper analyzes parallel- and sequential prototyping in new product development, modeling how time-to-market, prototyping mode and cost impact expected profit. We develop optimal policies for parallel and sequential prototyping, show how the parameters and form of the profit distribution drive the optimal policy, and evaluate the real option to abandon launch. We also specify a hybrid policy that optimizes time and resource tradeoffs. We find that the shape of the upper tail of the profit distribution, whether bounded, exponential, or fat-tailed, is a key driver. In short, the degree of uncertainty about upside profit potential determines the optimal level of experimentation and degree of parallelism.

The remainder of the paper proceeds as follows. In section 2 we study optimal parallel prototyping. Section 3 develops optimal hybrid policies that include parallel *and* sequential experiments. Section 4 concludes with a discussion of the results and managerial implications.

2. Parallel prototyping

2.1. The Model [1]

Assume the profit resulting from a single prototype is a random variable, X, with cumulative distribution function $F(x)$ and a continuous probability density function $f(x)$. X includes the net present value of incremental gross profits resulting from launching the new product, net of all costs other than the prototyping costs themselves. Hence, X captures the uncertainty about market risk, technology risk, and management risk. The parameters of the profit distribution derive from prior research on the market, technology, and competition facing the firm. This front-end process employs methods such as voice of the customer, target costing, lead user analysis, benchmarking, conjoint analysis, quality function deployment, parametric analysis and the like (c.f. Dahan [1997]).

[1] Model notation is summarized in Appendix A.

Profit outcomes for a firm developing n prototypes are modeled as independent draws from $F(x)$. Actually, this assumption can be relaxed. If outcomes are statistically *dependent* due to common additive factors, then it is sufficient to assume that, conditioning on the common factors, the *residual* profits are independent. In that case the profit distribution shifts, but the optimal parallel policy (ignoring the option to abandon) remains the same.

Prototypes cost c each and deducting the total cost of prototyping, $n \times c$, from gross profit determines net profit. The number of parallel prototypes, n, is treated as a continuous decision variable (even though it would be integer-valued in actuality).

2.2. Optimal Parallel Prototyping

We focus on pure parallel prototyping within a single period. When a tight development window limits the firm to one prototyping cycle, a single period model of pure parallel search fits. The firm decides on n^*, the optimal number of prototypes to build and test so as to maximize expected profit. After the profit outcomes for each of the n prototypes is revealed, the firm launches the one with the highest profit as its new product.

The cumulative distribution function of the maximum of n independent prototypes is $[F(x)]^n$, and the corresponding density function is $n \cdot f(x) \cdot [F(x)]^{n-1}$. Given that the prototype with the maximum observed profit is launched, the expected profit from building and testing n prototypes, is given by

$$E[p_n] = n \int_{-\infty}^{\infty} x \cdot [F(x)]^{n-1} \cdot f(x) dx - c \cdot n . \tag{1}$$

When the objective function (1) is strictly concave in n, the globally optimal number of prototypes, n^*, is determined by differentiating $E[p_n]$ with respect to n and solving the first order condition,

$$(2) \qquad \int_{-\infty}^{\infty}[n \cdot \ln F(x)+1] \cdot x \cdot [F(x)]^{n-1} \cdot f(x)dx = c \,,$$

which equates the marginal benefit of the n^{th} draw with its marginal cost. We note that the solution to equation (2) depends on the ratio between c and the scale of X, hence n^* does not change when X and c are both scaled by a constant.

The optimality condition implies:

$$(3) \qquad \frac{\partial E[p_{n^*}(n^*(c),c)]}{\partial c} = \frac{\partial E[p_{n^*}(n^*(c))]}{\partial n^*} \cdot \frac{\partial n^*(c)}{\partial c} + \frac{\partial E[p_{n^*}(c)]}{\partial c} = -n^*,$$

since $\frac{\partial E[p_{n^*}(n^*(c))]}{\partial n} = 0$ at n^*. Since $n^* > 0$ for any prototyping experiment worth running, $\frac{\partial E[p_{n^*}]}{\partial c} < 0$, so equation (3) demonstrates that a (small) reduction in unit prototyping cost impacts profit by n^*, the number of prototypes. That is, if the cost for each prototype is cut by one dollar, and if n^* remains (approximately) the same, the overall cost reduction will be n^*.

2.2.1. Real Option

When the option to abandon product launch is not available, equations (1) and (2) apply. But if the profit distribution, $F(x)$, allows negative outcomes (i.e., losses), then even the best of n^* draws may be negative, an event that occurs with probability $F(0)^{n^*}$. In that instance, the firm benefits from having the option to abandon product launch. The option to abandon may not be available if: (a) the firm is contractually obligated to deliver, (b) the product is a required component of a larger system, or (c) the firm pre-announced launch, thus creating an implicit contract with its customers. Also, personal incentives on the part of individual decision-makers may preclude the abandonment option from being exercised, even when that choice might be optimal for the firm. In cases where potential losses are small or unlikely, the option to abandon is virtually worthless, in which case early commitment to launch may provide a competitive advantage.

When the option to abandon is available, the expected profit from building and testing n prototypes improves compared with the no-option case (i.e. the option value is non-negative). The expected profit from n prototypes given the option to abandon, $E[p_n^{option}]$, is given by:

$$(4) \qquad E[p_n^{option}] = E[\max(0, X_1, ..., X_{n^*})] - c \cdot n = n \cdot \int_0^{\infty} x \cdot [F(x)]^{n-1} \cdot f(x)dx - c \cdot n .$$

The value of the abandonment option, given that the firm develops n prototypes, is given by the difference between (1) and (4),

$$(5) \qquad E[p_n^{option}] - E[p_n] = -\int_{-\infty}^{0} x \cdot n \cdot [F(x)]^{n-1} \cdot f(x)dx ,$$

which is the mean minimum loss, given that all n outcomes are negative, weighted by the corresponding probability of that event. In other words, the option value is the probability-weighted loss that is being avoided since the firm is not forced to launch.

Clearly, the option value at the optimum requires a comparison of $E[p_{n^{**}}^{option}]$ and $E[p_{n^*}]$, where n^{**} maximizes (4). The difference between the two is higher than in expression (5) with $n = n^*$ since n^* maximizes (1), but is only a feasible solution to (4). Hence, $E(p_{n^*}) \leq E(p_{n^*}^{option}) \leq E(p_{n^{**}}^{option})$.

2.3. Product Design and Extreme Value Theory

This section applies the statistical theory of extreme values to product development. *F(x)*, which characterizes prototype profit uncertainty, is determined by the process of concept generation depicted in the first two boxes of Figure 2. As described in chapter 5 of Ulrich and Eppinger (1995, p.78), "an effective development team will generate hundreds of concepts, of which 5 to 20 will merit serious consideration during the concept selection activity." Hence, we model *F(x)* as a distribution over product concepts, each of which is the maximum of a larger subset of product possibilities. In addition, as shown in the previous section, the value of the prototype chosen for launch is the maximum of the sample of prototype outcomes. When the maximum is taken over a large number of i.i.d. random variables, its asymptotic distribution is given by extreme value theory (Gumbel (1958), Galambos (1978)). The following theorem summarizes the pertinent results.

Theorem 1: [Galambos 1978]: Let *H(x)* be a distribution function from which *m* independent draws are taken. Then $\lim_{m\to\infty} H^m(x)$, the limiting distribution of the maximum of *m* draws from *H(x)*, converges to one of three distributions (with properly chosen a_m and b_m) or to none at all:

$$(6)\quad F_I^m(x) = e^{-\left(\frac{x-x_0}{b_m}\right)^{-a}} \quad \text{iff} \quad \lim_{t\to\infty} \frac{1-H(tx)}{1-H(t)} = x^{-a} \text{ for some } a \in (0,\infty),$$

$$(7)\quad F_{II}^m(x) = e^{-\left(\frac{x_0-x}{b_m}\right)^{a}} \quad \text{iff} \quad \lim_{t\to 0} \frac{1-H(x_0-tx)}{1-H(x_0-t)} = x^{a} \text{ for some } a \in (0,\infty), \text{ or}$$

$$(8)\quad F_{III}^{m}(x) = e^{-e^{-\frac{x-a_m}{b_m}}} \quad \text{iff} \quad \lim \frac{1-H(t+xR(t))}{1-H(t)} = e^{-x} \quad \text{where} \quad R(t) = \frac{1}{1-H(t)}\int_t^{\infty} 1-H(x)dx\,.$$

We assume that *H(x)*, the underlying profit distribution for the universe of possible design ideas, satisfies one of the limits in (6) - (8). Since product ideas that are good enough to make it onto Figure 2's "short list" are each a maximum from a large sample drawn from *H(x)*, *F(x)* takes the form of one of the three extreme-value distributions.

Appendix B summarizes the three extreme value distributions, provides their means and variances, and illustrates their unique property of closedness under maximization, which states that the highest of n draws from $F_i(x),\ i = I, II, III$ is also distributed $F_i(x)$ with modified parameters. Appendix C summarizes the appropriate limit tests and distribution parameters connecting a given underlying distribution to its limiting extreme value distribution.

The three extreme value distributions can be unified under a single, continuous model,

$$(9)\qquad F(x) = e^{-\left(1+\frac{x}{a}\right)^{-a}} \qquad \text{[von Mises 1936]},$$

where the distribution is Frechet if $a > 0$, Weibull if $a < 0$, and Gumbel as $|a| \to \infty$. This last fact suggests that for sufficiently large $|a|$, the optimal parallel prototyping results for the Gumbel distribution apply broadly. The parameter a captures essential information about the upper tail of the profit distribution. Our analysis reveals that lower absolute values of a lead to widely different optimal policies.

The extreme-value distributions, normalized to zero mean and unit variance, are shown in Figure 3 and interpreted in the context of new product development below. We note the pronounced variation in the upper quantiles of three normalized distributions, as exemplified by the 99th-percentile stars in Figure 3.

Figure 3: Densities for the three extreme value distributions (m = 0, s^2 = 1)

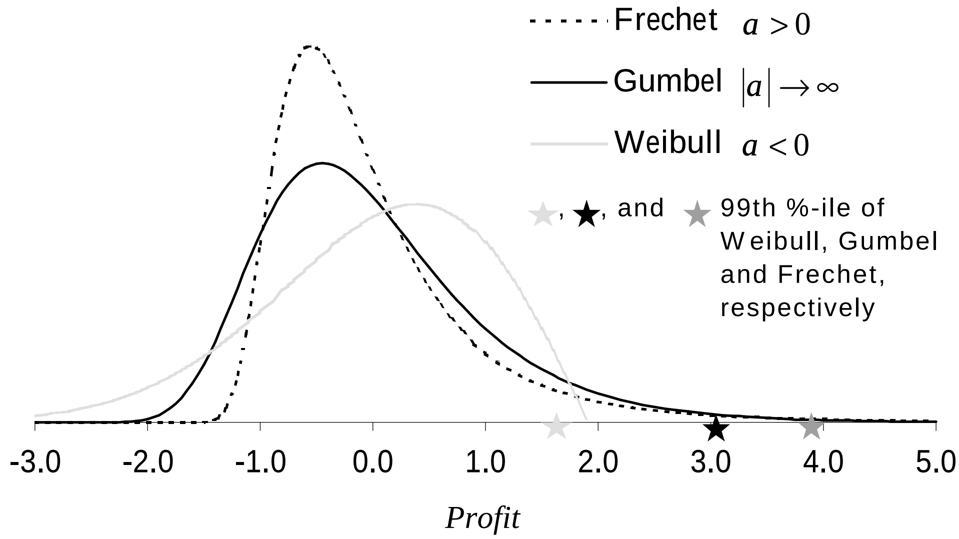

- *Weibull*: Some firms face predictably finite bounds on the upside profit potential of a new product due to limited market potential, price ceilings, or fixed price contracts. Such might be the case for a product that serves a small market, upgrades an existing user base, conforms to a fixed-price contract, or is capacity-constrained. When the gross profit is *upper-bounded*, the Weibull distribution applies. The gross profits from each prototype are distributed over an interval $[-\infty, x_0]$, where $x_0 > 0$.

- *Gumbel*: In many industries, there are no specific limits on the gross profit potential from a new product, but profit outcomes outside of a central range are extremely unlikely. Established products such as automobiles, food staples or commodities are not narrowly constrained by production capacity or market potential limits, but nevertheless tend towards somewhat predictable profit performance. When gross profit is *unbounded* from above, but with steeply declining probability density, a distribution with exponential tails is appropriate. The Gumbel distribution is the asymptotic distribution for the maximum of multiple draws from exponential-tailed distributions such as the normal.

- *Frechet*: Consider a product category with great upside uncertainty such as pharmaceuticals or new mass-market consumer durables. In such cases, products may become "mega-hits," accounting for the vast majority of the firm's profits. This may be due to network externality and dominant design effects, resulting in random variables that may be highly correlated in their effect on profit. When the gross profit distribution has a fat tail (e.g. when $F(x)$ declines as x^{-a}), the Frechet distribution applies with higher values of a denoting "thinner" upper tails. The gross profits from each prototype are then distributed over an interval $[x_0, \infty]$. As shown in Appendix B the Frechet distribution has infinite mean when $0 < a \leq 1$. Hence, we assume throughout that $a > 1$.

We next evaluate (1)-(5) and obtain closed form solutions for $n*$ for the Frechet, Weibull and Gumbel distributions.

2.4. Optimal Parallel Prototyping for the Three Extreme Value Distributions

2.4.1. Frechet Distribution

For the Frechet distribution, equation (1) becomes

$$E[p_n] = \frac{na}{b} \int_{x_0}^{\infty} x \cdot \left(\frac{x - x_0}{b} \right)^{-a-1} e^{-\left(\frac{x-x_0}{b} \right)^{-a}} dx - n \cdot c, \qquad (10)$$

which is maximized when

(11) $$n^* = \left[\frac{b}{ca}\Gamma\left(\frac{a-1}{a}\right)\right]^{\frac{a}{a-1}}.$$

Since $\frac{\partial^2 E[p_n]}{\partial n^2} = -b \cdot n^{\frac{1-2a}{a}} \cdot \left(\frac{a-1}{a^2}\right) \cdot \Gamma\left(\frac{a-1}{a}\right) < 0$ for all $a > 1$, $E[p_n]$ is strictly concave and n^* is globally optimal. As mentioned in the context of equation (2), it is the ratio of profit scale, b, to experimentation cost, c, that drives n^*. Thus, scaling b and c by the same constant in equation (11) leaves n^* unchanged. Also, $\frac{\partial n^*}{\partial a} < 0$ since higher values of a imply "thinner" tails.

We note that n^* is independent of the lower bound shift parameter, x_0 (and therefore of the mean) of the distribution. This is due to the fact that the benefit from extra experiments comes from incremental, relative improvements rather than absolute outcomes. The $n^{*\text{th}}$ experiment has an expected marginal benefit exactly equal to the cost of that experiment, c.

As the cost per experiment, c, declines, profit increases, $\frac{\P\, E[p_{n^*}]}{\P\, c} = -n^* < 0$, consistent with (3). As c declines, the optimal number of prototypes increases, $\frac{\partial n^*}{\partial c} < 0$.

Deriving the objective function as in (4) with the option to abandon yields:

(12) $$E\left[p_n^{opt}\right] = x_0 + bn^{1/a}\Gamma\left(\frac{a-1}{a}\right) - nc - x_0 e^{-n\left(-x_0/b\right)^{-a}} - bn^{1/a}\Gamma\left(\frac{a-1}{a}, n \cdot \left(\frac{-x_0}{b}\right)^{-a}\right).$$

Thus, the option to abandon is only relevant when the lower bound in the Frechet distribution is negative ($x_0 < 0$). Given that n prototypes are built, and subtracting

(10) from (12), the value of the option to abandon is $-x_0 e^{-n\left(-x_0/b\right)^{-a}} - bn^{1/a}\Gamma\left(\frac{a-1}{a}, n\cdot\left(\frac{-x_0}{b}\right)^{-a}\right)$, which can be interpreted as follows. The expression is the probability weighted loss that would be avoided when the option to abandon would be exercised. In particular, as the downside risk increases (i.e., as x_0 becomes more negative and losses more likely), both terms of the option value increase.

2.4.2. Weibull Distribution

For the Weibull distribution, equation (1) becomes

$$(13) \qquad E[p_n] = n\int_{x_0}^{\infty} x \cdot \frac{a}{b}\left(\frac{x_0 - x}{b}\right)^{-a-1} e^{-\left(\frac{x_0-x}{b}\right)^{-a}} dx - c\cdot n,$$

which is maximized when

$$(14) \qquad n^* = \left[\frac{b}{ca}\Gamma\left(\frac{a+1}{a}\right)\right]^{\frac{a}{a+1}}.$$

Since $\frac{\partial^2 E[p_n]}{\partial n^2} = -\frac{1+a}{a^2} bn^{-\frac{1+2a}{a}}\Gamma\left(\frac{a+1}{a}\right) < 0$, $E[p_n]$ is strictly concave and n^* is globally optimal. Again we note that scaling b and c by a constant leaves n^* unchanged and that n^* is independent of the upper bound of the distribution, x_0, since in the no option case the decision on n compares relative, incremental improvements in expected outcomes to the marginal cost of those improvements.

As the cost per experiment, c, declines, the optimal number of prototypes increases, $\frac{\partial n^*}{\partial c} < 0$ and, once again, $\frac{¶ E[p_{n^*}]}{¶ c} = -n^* < 0$, consistent with (3).

Deriving the objective function with the option to abandon as in (4) yields:

$$E\left[p_n^{opt}\right] = x_0 - bn^{-1/a}\Gamma\left(\frac{a+1}{a}\right) - nc - x_0 e^{-n\left(x_0/b\right)^a} + bn^{-1/a}\Gamma\left(\frac{a+1}{a}, n\cdot\left(\frac{x_0}{b}\right)^a\right). \tag{15}$$

First, we note that the Weibull distribution is only relevant when the upper bound is positive, that is, when $x_0 > 0$, since otherwise profits would not be possible and no experimentation would take place. Given that n prototypes are built, and subtracting (13) from (15), the value is $-x_0 e^{-n\left(x_0/b\right)^a} + bn^{-1/a}\Gamma\left(\frac{a+1}{a}, n\cdot\left(\frac{x_0}{b}\right)^a\right)$, which is the probability weighted loss that would be avoided when the option to abandon is exercised. In particular, as the upper bound of the Weibull, x_0, decreases towards zero, both terms of the option value increase.

2.4.3. Gumbel Distribution

Employing the Gumbel's closure under maximization, the expected maximum profit from n draws is given by

$$E\left[p_n\right] = a + b\ln n + bg - cn. \tag{16}$$

Since $\frac{\P\, E^2\left[p_n\right]}{\P\, n^2} = -\frac{b}{n^2} < 0$, $E[p_n]$ is strictly concave and n^* is a global maximum. The first order condition leads to the remarkably simple expression

$$n^* = \frac{b}{c}. \tag{17}$$

Equation (17) again reveals that the optimal number of prototypes is independent of the *mean* of the underlying profit distribution since the shift parameter a does not appear in the expression. Rather, the result depends only on the ratio between the profit scale and the cost per experiment. This simple result is intuitively appealing in that both declining prototyping costs and greater profit uncertainty increase prototyping activity.

As the cost per experiment, *c*, declines, the optimal number of prototypes increases, $\frac{\P\, n^*}{\P\, c} = -\frac{b}{c^2} < 0$. As *c* declines, profit increases, $\frac{\P\, E[p_{n^*}]}{\P\, c} = -\frac{b}{c} < 0$, consistent with (3).

Deriving the objective function as before, but retaining the option to abandon the project if even the best of n prototypes generates a loss, yields:

$$E\left[p_n^{opt}\right] = a + b\ln n + bg - cn + \left[-b \cdot \mathrm{Ei}\left(-ne^{a/b}\right)\right], \qquad (18)$$

where Ei(y) is the exponential integral function, $-\mathrm{Ei}(y) = \int_{-\infty}^{y} \frac{e^{-y}}{y} dy$. Given that n prototypes are built, the value of the option to abandon is $\left[-b \cdot \mathrm{Ei}\left(-ne^{a/b}\right)\right]$, which can be interpreted as follows. For negative values of y, -Ei(y) is always positive and decreases in $|y|$. Thus, as $ne^{a/b}$ increases, $\left[-b \cdot \mathrm{Ei}\left(-ne^{a/b}\right)\right]$ stays positive, but decreases in magnitude. Indeed, the option to abandon has the greatest value when only a few experiments are run (low n), when uncertainty is very high (high *b*), or when mean profit is low (low *a*). When many experiments are run (high n), and high profits (high *a*) are relatively certain (low *b*), the option value approaches zero, as one might expect.

The comparative statics results when the option to abandon exists fit our intuition. As profit variability, *b*, increases, so does the optimal number of prototypes, since the possibility of at least one very high outcome increases and we must remember that only the highest outcome matters. As *a*, the shift parameter of the distribution, is increased, the optimal number of prototypes *increases*. The reason is that when *a* increases, downside costs are less likely, so the option value $\left[-b \cdot \mathrm{Ei}\left(-ne^{a/b}\right)\right]$ decreases. In the no option case, upward or downward shifts in the distribution (i.e., changes to *a*) have no effect on n^*.

Figure 4 depicts expected profits for Frechet, Gumbel and Weibull profit distributions which have been normalized to zero mean and unit variance. The two cases of experimentation costs $c = 0.001$ and $c = 0.1$ are depicted as the tail shape parameter, a, increases towards infinity.

Figure 4

Effect of Tail Shape Parameter a on Expected Profit

c=0.1 or 0.001, $F_i(x)$ normalized to m=0, s^2=1

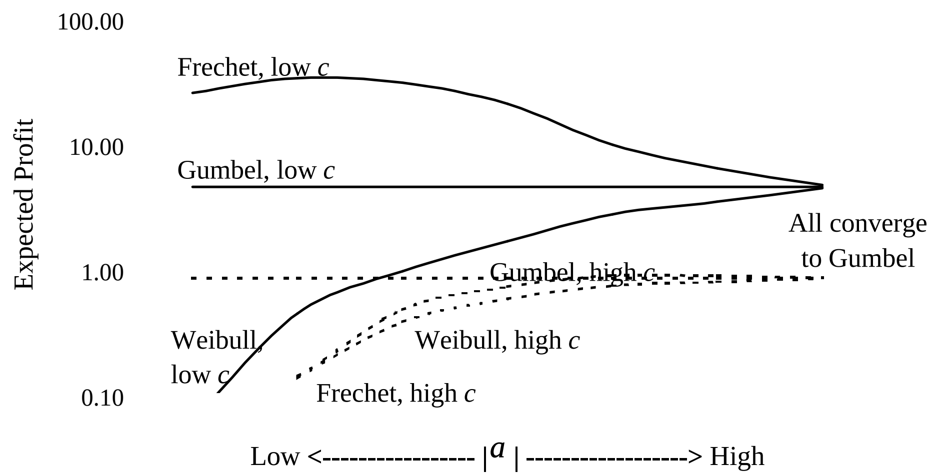

As the absolute value of the tail shape parameter, a, increases, we note that the Frechet and Weibull distributions converge to the Gumbel. Thus, the optimal number of parallel prototypes and the expected profits resulting from building them also converge as $|a|$ increases. For lower values of $|a|$, however, the three distributions diverge greatly in their profit implications and in their optimal number of experiments, particularly when experimentation costs are low. Holding mean and variance constant, lower $|a|$ values imply fatter tails in the case of the Frechet distribution and tighter upper bounds in the Weibull case. Thus, the upper-tail-shape of the profit distribution, as parameterized by a, plays a pivotal role in determining the optimal experimentation policy and the expected profit that results from following that policy.

2.5. Total Prototyping Spending as Unit Prototyping Costs Decline

How do declining prototyping costs affect total spending for firms conducting experiments in parallel?

Figure 5: Total Prototype Expenditures as Prototype Costs Decrease

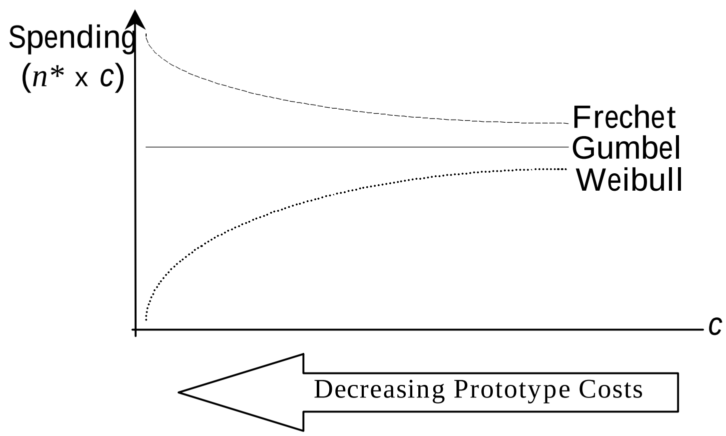

In the Frechet case, reducing the unit prototyping cost *increases* optimal total prototyping expenditures, $\frac{\partial (n^* \times c)}{\partial c} < 0$ when $a > 1$[2], in a convex fashion, i.e. $\frac{\partial^2 (n^* \times c)}{\partial c^2} > 0$. In other words, the firm's demand for prototyping is elastic when profits follow a Frechet distribution. Thus, when profits are highly uncertain on the upside, the firm should exploit lower-cost prototyping technologies by increasing not only the

[2] Recall $a > 1$ is the case of interest for the Frechet since it is required for the mean to remain finite.

number of parallel experiments, but also the total amount spent on those experiments. This behavior contrasts with total expenditures for the Weibull and Gumbel profit distributions as depicted in Figure 5.

In the Weibull case, the induced demand for prototyping is inelastic as declining unit prototyping costs concavely reduce optimal total expenditures, that is, $\frac{¶\,(n^* \times c)}{¶\,c} > 0$ and $\frac{¶\,^2(n^* \times c)}{¶\,c^2} < 0$. Intuitively, since Weibull-regime firms are constrained in their upside profit potential, the marginal benefits of additional experiments are small, hence total spending declines as prototyping costs drop.

In the Gumbel case, reducing the unit prototyping cost has no effect on total spending, ($\frac{¶\,(n^* \times c)}{¶\,c} = 0$), since total spending stays constant at *b*, the scale parameter of the distribution. Thus, the induced demand for prototyping has unit elasticity in the Gumbel case. To the extent that the Gumbel distribution captures the underlying profit behavior of many product development efforts, the unit elasticity result explains the relative stability of total prototyping spending as a percentage of firm profits even in the context of rapidly declining unit prototyping costs.

We can quantify the opportunity cost of sub-optimal spending by evaluating the constrained optimization problem: $\max_n \left\{ n \cdot \int_{-\infty}^{\infty} x \cdot [F(x)]^{n-1} \cdot f(x)dx - c \cdot n \right\}$ subject to $c \cdot n \leq M$, where M is the prototyping budget constraint. The shadow price on prototyping budget is $l = \frac{1}{c} \int_{-\infty}^{\infty} \left[\frac{M}{c} \cdot \ln F(x) + 1 \right] \cdot x \cdot [F(x)]^{\frac{M}{c}-1} \cdot f(x)dx - 1$. For example, equation (17) gives total optimal spending for the Gumbel case as $n^* \times c = b$. When the budget is an amount

$M < b$, so that only $M/c < n^*$ prototypes can be built, the shadow price on an extra prototyping dollar takes the simple form $l_{Gumbel} = \frac{b}{M} - 1$. When $M \geq b$, the constraint does not bind, but when $M < b$, it is optimal to "rent" additional prototyping resources. Similar analyses for the Frechet and Weibull distributions lead to $l_{Frechet} = \frac{b}{ca}\Gamma\left(\frac{a-1}{a}\right)\cdot\left(\frac{c}{M}\right)^{\frac{a-1}{a}} - 1$ and $l_{Weibull} = \frac{b}{ca}\Gamma\left(\frac{a+1}{a}\right)\cdot\left(\frac{c}{M}\right)^{\frac{a+1}{a}} - 1$, respectively. If only one prototype is budgeted $(M = c)$, then $l_{Gumbel} = \frac{b}{c} - 1$, $l_{Frechet} = \frac{b}{ca}\Gamma\left(\frac{a-1}{a}\right) - 1$, and $l_{Weibull} = \frac{b}{ca}\Gamma\left(\frac{a+1}{a}\right) - 1$.

3. Hybrid Parallel/Sequential Policies

While most firms are organized to take advantage of pure sequential or pure parallel prototyping, but not both, some "ambidextrous" firms may be able to conduct a hybrid sequential/parallel policy. Such firms conduct smaller parallel efforts, observe the best outcome from each effort, and iterate until a good enough result is observed. By balancing the relationship between cost, expected profit and the need for speed-to-market, ambidextrous firms can globally optimize their net expected profit.

In this section we relax the constraint that all prototypes must be built in a single period. We assume an infinite time horizon with discount factor d per period, $0 < d \leq 1$ (periods are measured relative to the end of period one). Any number of i.i.d. prototyping experiments can be run in each period. We show that for sufficiently low d, the optimal policy can be characterized as a hybrid sequential/parallel policy in which m^* parallel prototypes are built in each period until a "good enough" result is achieved.

If d is sufficiently close to 1, that is if delays to market are not very costly, then a pure sequential policy will be optimal. We conclude by comparing the expected profits of the optimal pure parallel, pure sequential, and hybrid policies.

3.1. Purely Sequential Experimentation

As a basis for comparison against pure parallel and hybrid parallel/sequential processes, we use the result in Weitzman [1979] for the optimal pure sequential search policy, referred to as *Pandora's rule* since it consists of opening one box at a time (building a single prototype) with unpredictable contents (observing a stochastic outcome). Each possible prototype is parameterized by the cost of building it and the probability distribution of possible rewards. A reservation price, *z*, is assigned to each experiment and is the solution to the following equation:

$$z = -c + \int_{z}^{\infty} x \cdot f(x)dx + d \cdot z \cdot \int_{-\infty}^{z} f(x)dx \,, \tag{19}$$

which can also be written as:

$$c = \int_{z}^{\infty} (x - z) f(x)dx + (1 - d) \cdot z \cdot F(z) \,, \tag{20}$$

where *c* is the cost of the prototype, $F(x)$ the density function of the profit distribution, and d the discount factor per period. Equation (19) reveals that the reservation price equals the net expected value of running sequential experiments until realizing the first outcome greater than or equal to the reservation price. That is, *z* is the expected value of following a sequential experimentation policy in which the cutoff value for stopping is *z* itself. From equation (20) we see that *z* increases monotonically in d (intuitively, if future dollars are worth more, then the cutoff for stopping should be higher since we are more willing to wait for the reward) and decreases monotonically in *c*.

Weitzman proves that the optimal pure sequential search policy is to open the box with the highest reservation price, observe the stochastic outcome, and stop if the outcome is higher than the reservation prices for all remaining boxes. When replacement is permitted, for example when an unlimited number of the same box may be opened, Weitzman's optimal policy suggests opening just one type of box, the one with the highest reservation price, until exceeding that reservation price.

When delay is of no consequence, $d = 1$ and there is no economic advantage to building prototypes in parallel. The optimal policy is then purely sequential with one prototype built at a time order according to Weitzman's Pandora's rule. Figure 6 illustrates the benefits of pure sequential experimentation when there is no discounting.

Figure 6

Effect of tail shape a on Profit of Sequential vs. Parallel
Assuming no Discounting (d=1) and c=0.1 or 0.001

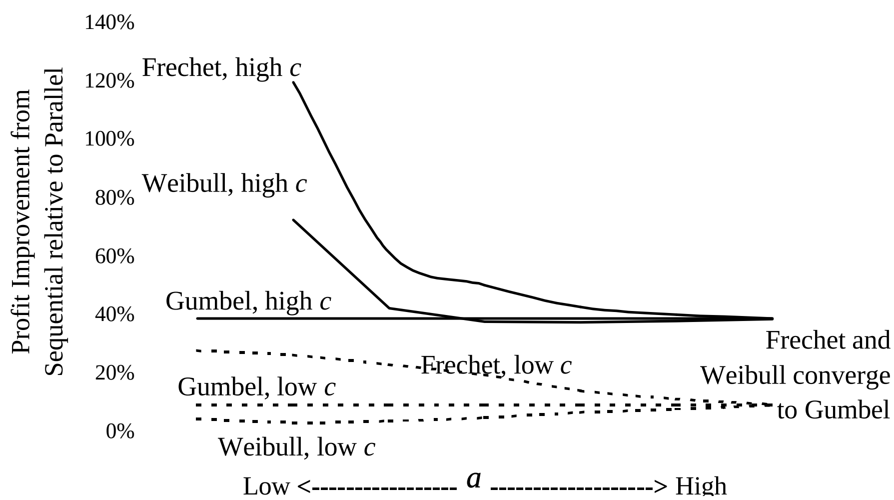

In general, higher prototyping costs are seen to enhance the relative profit advantage of sequential experimentation for all distributions since parallel experimentation grows too costly. Fatter tails in the Frechet distribution (lower a's) also favor sequential experimentation as the firm sets ever higher expectations for what level of profit performance is "good enough."

When discounting exists, as would be the case when time-to-market is relevant, the relative profit improvement due to sequential experimentation declines with the discount factor. In fact, with a discount factor at or below 0.9 per period, the profit improvements depicted in Figure 6 disappear in all but the most extreme cases (e.g. $c > 0.1$ and $a < 2$).

For $d < 1$, the following proposition shows that there exists a discount factor, d_{Switch}, for which the expected value of following a pure parallel policy actually exceeds that of following a pure sequential policy.

Proposition 1: Given experiment cost c and profit distribution $F(x)$, let $z(d)$ solve equation (20). Let $E[p_{n^*}]$ be the expected value of following the optimal pure parallel policy, given c and $F(x)$. Assume $E[p_{n^*}] > E[p_1]$, the expected profit from a one-shot policy. Then there exists d_{Switch} such that: (1) $z(d_{Switch}) = E[p_{n^*}]$, (2) $z(d) < E[p_{n^*}]$ for $d < d_{Switch}$, and (3) $z(d) > E[p_{n^*}]$ for $d > d_{Switch}$.

Proof: When $d = 0$, $z(0) = E[p_1] < E[p_{n^*}]$, since rewards from future periods are completely discounted. When $d = 1$, $z(1) \geq E[p_{n^*}]$ since, when there is no discounting, running n^* or fewer sequential experiments is at least as good as running n^* at the same time. Equation (20) gives us that $z(d)$ increases monotonically in d and since $0 \leq d \leq 1$, this completes the proof.

3.2. Optimal hybrid Policy

Next we consider a hybrid parallel/sequential policy in an infinite-horizon setting. In order to determine the optimal hybrid policy, we note that m i.i.d. parallel experiments can be viewed as a *single*, composite experiment with prototyping cost $m \cdot c$, distribution function $F_m(x) = [F(x)]^m$, and density $f_m(x) = m \cdot f(x) \cdot [F(x)]^{m-1}$. Thus, the hybrid problem becomes a special case of the optimal pure sequential problem, where the choice of the number of parallel experiments within a period is recast as a choice from alternative experiments, parameterized by m. The corresponding reservation price, z_m, solves

$$(21)\qquad z_m = -m \cdot c + \int_{z_m}^{\infty} x \cdot m \cdot f(x) \cdot [F(x)]^{m-1} dx + d \cdot z_m \cdot [F(z_m)]^m .$$

By maximizing m, we obtain the optimal composite experiment (consisting of m^* parallel prototypes) with the highest net expected value. Since the experiment with the highest reservation price, z_{m^*}, should be run first, it follows that the optimal hybrid policy in the infinite horizon problem is to run m^* experiments in parallel in each period until a result greater than z_{m^*} is observed. The expected profits resulting from the optimal hybrid policy is z_{m^*}, which can be compared with the expected profits from pure sequential and pure parallel policies as depicted in Figure 7.

Figure 7: Comparative Performance as a function of d

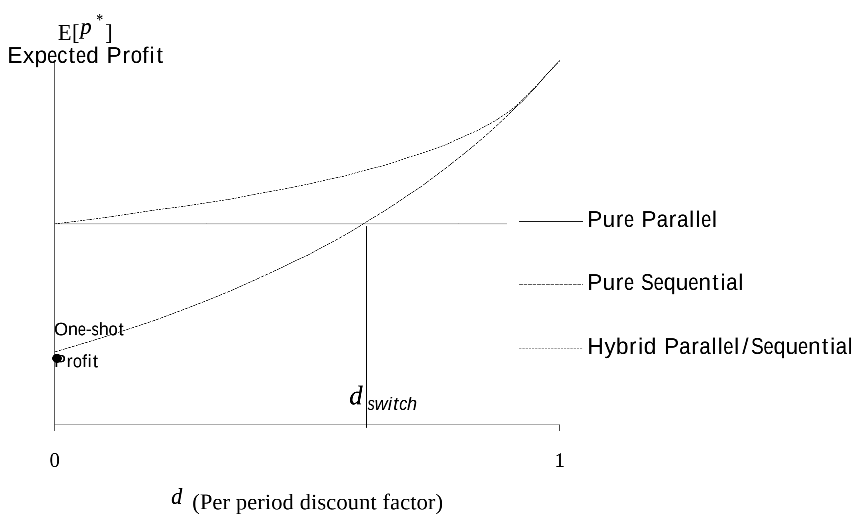

In fact, as Proposition 2 shows, it is clear that the optimal hybrid policy strictly dominates either pure parallel policy over an interior range of discount factors d.

Proposition 2: Given experiment cost c and profit distribution $F(x)$, let $z_{m^*}(d)$, the expected value of following the optimal hybrid policy, solve equation (21). Let $z(d)$, the expected value of following a pure sequential policy, solve equation (20). Finally, let $E[p_{n^*}]$, the expected value of following the optimal pure parallel policy, maximize equation (1) and exceed the expected value of a one-shot policy. Then there exists a range (d_{Low}, d_{High}) such that $z_{m^*}(d) > \max(z(d), E[p_{n^*}])$ for all $d \in (d_{Low}, d_{High})$. Further, $d_{Switch} \in (d_{Low}, d_{High})$.

Proof: By definition, $0 \le d \le 1$. When $d = 0$, rewards from future periods are completely discounted so $m^* = n^*$ and $z_{m^*}(0) = E[p_{n^*}]$, since the problem reduces to a single-period optimization. When $d = 1$, $z_{m^*}(1) = z(1)$ and $m^* = 1$ since only individual, sequential experiments will be run when there is no discounting. Equation (20) gives us that $z(d)$ and $z_{m^*}(d)$ both increase monotonically in d (letting $F(x) = F^{m^*}(x)$ and $c = m^* \cdot c$ in case of $z_{m^*}(d)$). Since $z(d)$ increases monotonically from $z(0)$ to $E[p_{n^*}]$ over the range $d \in [0, d_{Switch}]$ from Proposition 1, and $z_{m^*}(d)$ increases monotonically over the same range from a starting point of $z_{m^*}(0) = E[p_{n^*}] > E[p_1] = z(0)$, then clearly the proposition holds over the range $(0, d_{Switch}]$ for some $d_{Low} \in (0, d_{Switch})$. Likewise, over the range $d \in [d_{Switch}, 1]$, both $z_{m^*}(d)$ and $z(d)$ monotonically increase until they are equal at $d = 1$, but $z_{m^*}(d)$ starts out at a higher value, so the proposition holds over the range $[d_{Switch}, 1)$ for some $d_{High} \in (d_{Switch}, 1)$. This completes the proof.

The optimal hybrid parallel/sequential cutoff values for Frechet, Weibull and Gumbel distributed prototyping experiments depend on *d*, the discount factor per period, in that lower values of *d* reduce the present value of profits from later design iterations in a sequential prototyping mode. Heavy discounting for time-to-market delays makes parallel prototyping a more attractive choice. Applying (19) to the cases of Frechet, Weibull and Gumbel distributed profit, but with *F(x)* replaced by the distribution function $F_m(x) = [F(x)]^m$, the cutoff values for the optimal hybrid policies satisfy the following equations:

(22) Hybrid Frechet:
$$z_F^H = -m \cdot c + d \cdot z_F^H e^{-m\left(\frac{z_F^H - x_0}{b}\right)^{-a}} + x_0 \left(1 - e^{-m\left(\frac{z_F^H - x_0}{b}\right)^{-a}}\right) + b \cdot m^{1/a}\left[\Gamma\left(1 - \frac{1}{a}\right) - \Gamma\left(1 - \frac{1}{a}, m \cdot \left(\frac{z_F^H - x_0}{b}\right)^{-a}\right)\right]$$

(23) Hybrid Weibull:
$$z_W^H = -m \cdot c + d \cdot z_W^H e^{-m\left(\frac{x_0 - z_W^H}{b}\right)^{a}} + x_0 \left(1 - e^{-m\left(\frac{x_0 - z_W^H}{b}\right)^{a}}\right) - b \cdot m^{-1/a}\left[\Gamma\left(1 + \frac{1}{a}\right) - \Gamma\left(1 + \frac{1}{a}, m \cdot \left(\frac{x_0 - z_W^H}{b}\right)^{a}\right)\right]$$

(24) Hybrid Gumbel:
$$z_G^H = -m \cdot c - (1-d) \cdot z_G^H \cdot e^{-m \cdot e^{-\left(\frac{z_G^H - a}{b}\right)}} + a + b \cdot g + b \ln m - b \cdot \mathrm{Ei}\left(-m \cdot e^{-\left(\frac{z_G^H - a}{b}\right)}\right)$$

Equations (22)-(24) do not have closed-form solutions, but are useful in deriving comparative statics results and can be solved numerically for the optimal policy. For example, it becomes clear that lowering the cost per prototype, c, increases the reservation price z, while lowering d, the per period discount factor, reduces z and increases m^*, the optimal number of experiments per cycle.

We can compare the performance of the optimal hybrid policy to that of either pure policy. Numeric results showing relative profit performance under the three distributions, all with zero mean and unit variance, are presented in Figure 8 for the case of low ($c = 0.001$) and high ($c = 0.1$) experimentation costs. A period-to-period discount factor of $d = 0.9$ is assumed, a period being defined as the time needed to complete one round of experimentation.

Figure 8

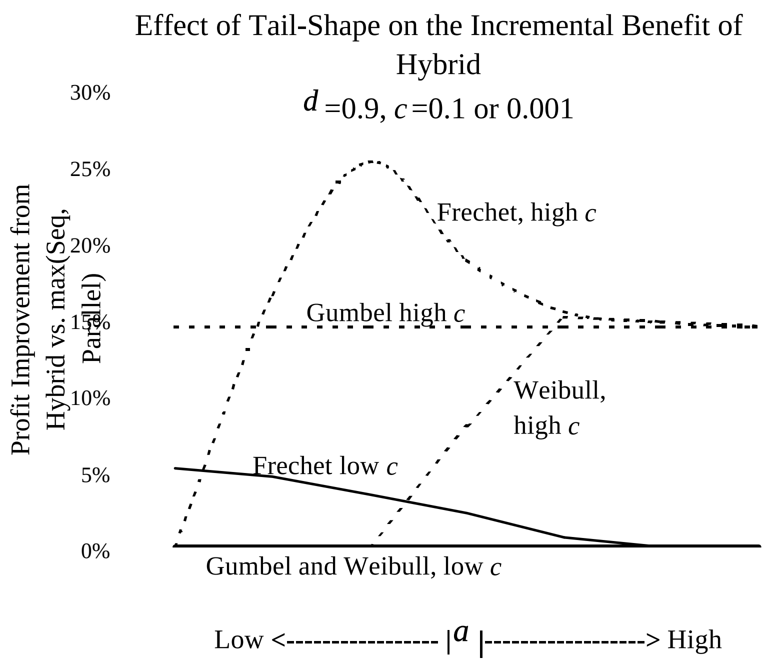

These results indicate that while the optimal hybrid policy dominates either pure policy as shown in Proposition 2, its incremental impact on expected profit only exceeds 15% when experimentation is costly or the profit distribution has a fat tail. Otherwise, when discounting is minimal and/or experiments costly, the pure sequential policy will perform almost as well as the hybrid. And when discounting is heavy and/or experimentation costs low, the pure parallel policy will perform as well as the hybrid.

4. Concluding Remarks

This paper investigates optimal policies of parallel and sequential prototyping from economic and probabilistic perspectives and quantifies the tradeoffs between development speed, development cost, and profit performance.

Time-to-market pressures, parameterized by a discount factor d, are shown to favor parallel over sequential prototyping when d is less than a unique cutoff discount factor, d_{switch}. Thus, the trends towards shorter product development cycles and accelerating industry clockspeeds (Mendelson and Pillai (1998)) favor parallel prototyping. We specify the optimal hybrid policy that dominates the profit performance of both pure policies by combining the speed advantage of parallelism with the development cost advantage of sequentialism.

Our model utilizes the statistical theory of extreme values to calculate optimal parallel prototyping policies. We demonstrate the Gumbel distribution's central role in characterizing profit uncertainty for product concepts that are selected as the best of many ideas considered. Our analysis leads to the remarkably simple result that for Gumbel-distributed profits, the optimal number of parallel prototypes is $n^* = b/c$, where b is the scale parameter for the distribution and c is the cost per prototype. Thus, the optimal number of prototypes depends only on the ratio between the profit distribution's scale parameter and the cost per prototype.

Parallel prototyping policies under the other extreme value distributions, i.e. those for which the tail-shape parameter $|a|$ is small, are also shown to depend on b/c. However, expected profit under low-$|a|$ regimes, even with the same mean and variance, diverges significantly from that under a Gumbel (infinite-$|a|$) regime. When a is negative, but small, the upside is bounded and both expected profits and optimal number of parallel experiments are dramatically lower. Conversely, when a is positive, but small, profits are fat-tailed and the expected profit and number of parallel experiments are dramatically higher. Thus, we show that the upper tail-shape of the profit distribution drives the number of parallel prototypes and the extent to which they are expected to pay off.

The effect of lower unit prototyping costs on total prototyping spending also depends on the profit distribution tail-shape parameter, a. Development costs remain stable when profits are Gumbel-distributed, with reductions in unit prototyping costs being exactly offset by an increased number of prototypes. But declining unit prototyping costs under a bounded profit distribution lead to reductions in total development spending; the firm "takes the money and runs". In contrast, declining unit prototyping costs under a fat-tailed profit distribution lead to increases in total development spending.

In short, parallel prototyping gains importance when: (a) time-to-market pressures increase (d declines), (b) upside profits are less certain (i.e., are distributed with fatter upper tails), or (c) experimentation costs decline.

Firms seeking to maximize new product profits should: (1) encourage creativity and experimentation as a way of fattening the upper tail of the profit distribution, (2) invest in cost efficient prototyping technologies to lower unit experimentation costs, and (3) remember that "time is money" when setting R&D capacities and budgets.

The firm may be able to enhance its profit distribution by supporting a culture of creativity, involving outside idea suppliers, and encouraging experimentation even if it ultimately leads to failure. Having design teams compete may also improve the profit distribution, but members of the losing team will require proper incentives to "fail forward" (Leonard-Barton 1995). Intel, for example, holds a big party to honor the "losers". Competition between design teams may improve results, speed up the design process, and create incentives to perform. Where internal competition might cause strife or where prototyping resources are limited, external suppliers can be hired to build and test multiple prototypes.

The trend towards lower unit prototype costs is fueled by investment in technologies such as virtual design, rapid prototyping, combinatorial methods and automated processing. Lower unit costs may also derive from economies of scale inherent in parallel prototyping itself, since some costs may be fixed. Combinatorial chemists can automatically test the performance of one hundred compounds on a test substrate just as they can ten. And consumers can respond to multiple designs on a web page as easily as they can to one, without significantly impacting the cost of market testing. Recent work by Dahan and Srinivasan (1998) demonstrates that virtual prototypes on the Web result in market share predictions that are nearly identical to those for costlier physical prototypes. We could soon witness a period of widespread virtual parallel prototyping and market testing.

Organizational structure also affects the suitability of one-shot, sequential and parallel modes of prototyping. Firms may be organized to better utilize one mode or another, but design teams that are ambidextrous, and can successfully implement hybrid policies of parallel *and* sequential prototyping, can gain a distinct advantage over their competitors.

Parallel prototypes may be actual products. Firms engaging in mass-customization may launch multiple permutations of a product, allowing customers to express their preferences directly through purchase. After observing real demand, the firm can focus on the most profitable designs. Seiko watches, Motorola pagers, and Dell personal computers embody this product-as-prototype approach. In fact, for some of its personal computers, according to Dell Vice President Stuart Smith, Dell has had only a single shipment.

In summary, the choice of prototyping mode profoundly impacts the profitability of products that the firm develops.

Appendix A: Model parameters, variables and notation

n	Number of prototypes to be built and tested; a decision variable.
$n*$	Optimal number of prototypes to build without the abandonment option
$n**$	Optimal number of prototypes when the option to abandon is available
$m*$	Optimal number of prototypes per period for a hybrid policy
c	Cost to build and test each prototype
X	Random variable for the gross profit from a single prototype
Z_n	Random variable for $\max(X_1, X_2, \cdots, X_n)$
$F(x)$	Cumulative distribution function for X
$F(x)$	Probability density function for X
b	Discount factor per period, $0 < b \leq 1$
b_{switch}	b producing equal expected profit for parallel and sequential policies
p_n	Random variable for the maximum net profit available after n draws, $p_n = Z_n - n \cdot c$
p_n^{option}	Random variable for Max ($0,\ p_n$), i.e., the maximum net profit from n draws, given the option to abandon
z_n	Reservation price used to optimally order sequential experiments; Solves $z_n = -n \cdot c + \int_{z_n}^{\infty} x \cdot n \cdot f(x) \cdot [F(x)]^{n-1} dx + b \cdot z_n \cdot [F(x)]^n$

Appendix B: Summary of Frechet, Weibull and Gumbel Distributions

Cumulative Distribution and Probability Density	Mean and Variance	Closedness Under Maximization (n draws)
Frechet (Type-I) $F_I(x) = e^{-\left(\frac{x-x_0}{b}\right)^{-a}}$ $f_{I(x_0,b,a)}(x) = \frac{a}{b}\left(\frac{x-x_0}{b}\right)^{-a-1} e^{-\left(\frac{x-x_0}{b}\right)^{-a}}$	$E[x] = x_0 + b\cdot\Gamma\left(\frac{a-1}{a}\right)$ where $\Gamma(t) = \int_0^\infty x^{t-1}e^{-t}dt$, the Gamma function $\mathrm{Var}[x] = b^2\Gamma\left(\frac{a-2}{a}\right) - \left[b\Gamma\left(\frac{a-1}{a}\right)\right]^2$	$\left[F_{I(x_0,b,a)}(x)\right]^n = e^{-n\left(\frac{x-x_0}{b}\right)^{-a}}$ $= e^{-\left(\frac{x-x_0}{bn^{\frac{1}{a}}}\right)^{-a}}$ $= F_{I(x_0,b\cdot n^{\frac{1}{a}},a)}(x)$
Weibull (Type-II) $F_{II}(x) = e^{-\left(\frac{x_0-x}{b}\right)^{a}}$ $f_{II(x_0,b,a)}(x) = \frac{a}{b}\left(\frac{a-x}{b}\right)^{a-1} e^{-\left(\frac{a-x}{b}\right)}$	$E[x] = x_0 - b\cdot\Gamma\left(\frac{a+1}{a}\right)$ $\mathrm{Var}[x] = b^2\Gamma\left(\frac{a+2}{a}\right) - \left[b\Gamma\left(\frac{a+1}{a}\right)\right]^2$	$\left[F_{II(x_0,b,a)}(x)\right]^n = e^{-n\left(\frac{x_0-x}{b}\right)^{a}}$ $= e^{-\left(\frac{x_0-x}{bn^{-\frac{1}{a}}}\right)^{a}}$ $= F_{II(x_0,b\cdot n^{-\frac{1}{a}},a)}(x)$
Gumbel (Type-III) $F_{III}(x) = e^{-e^{-\frac{x-a}{b}}}$ $f_{III(a,b)}(x) = \frac{1}{b}e^{-\left(\frac{x-a}{b}\right)}e^{-e^{-\left(\frac{x-a}{b}\right)}}$	$E[x] = a + b\cdot g$ where $g \cong 0.57722$ is Euler's constant $\mathrm{Var}[x] = b^2\frac{p^2}{6}$	$\left[F_{III(a,b)}(x)\right]^n = e^{-n\cdot e^{-\frac{x-a}{b}}}$ $= e^{-e^{-\frac{x-(a+b\ln n)}{b}}}$ $= F_{III(a+b\ln n,b)}(x)$

Appendix C: Summary of Extreme Value Theory Limiting Distributions

Required Limit for the Underlying Distribution	Limiting Distribution (parent examples)	a_m and b_m
$\lim_{t\to\infty}\frac{1-H(tx)}{1-H(t)}=x^{-a}$ Requires: $x_0 < x \le \infty$ and $a>0$ ($a>1$ for finite mean and $a>2$ for finite variance) Note: $\lim_{x\to\infty}\frac{xh(x)}{1-H(x)}=a$	Frechet (Type-I) $F_I^m(x)=e^{-\left(\frac{x-x_0}{b_m}\right)^{-a}}$ (e.g. Pareto parent such as $H(x)=1-\frac{1}{x^a}$)	$b_m=\inf\left\{x:1-H(x)\le\frac{1}{m}\right\}$ Note: $0<g\le 1$ results in $m=\infty$
$\lim_{t\to 0}\frac{1-H(x_0-tx)}{1-H(x_0-t)}=x^{a}$ Requires: $-\infty<x\le x_0$ and $a>0$	Weibull (Type-II) $F_{II}^m(x)=e^{-\left(\frac{x_0-x}{b_m}\right)^{a}}$ (e.g. Uniform parent such as $H(x)=x$)	$b_m=x_0-\inf\left\{x:1-H(x)\le\frac{1}{m}\right\}$
$\lim\frac{1-H(t+xR(t))}{1-H(t)}=e^{-x}$, where $R(t)=\frac{1}{1-F(t)}\int_t^\infty 1-F(x)dx$	Gumbel (Type-III) $F_{III}^m(x)=e^{-e^{-\frac{x-a_m}{b_m}}}$ (e.g. Exponential parent such as $H(x)=1-e^{-x}$)	$a_m=\inf\left\{x:1-F(x)\le\frac{1}{m}\right\}$, and $b_m=R(a_m)$ $=m\int_{a_m}^\infty 1-F(x)dx$

References

Baldwin, Carliss Y. and Kim B. Clark. "Modularity-in-Design: An Analysis Based on the Theory of Real Options." Working Paper. July, 1994. 46 pages.

Dahan, Ely. "Listening to the Customer." Teaching Note. 1998.

Dahan, Ely and V. Srinivasan. "The predictive Power of Internet-Based Product Concept Testing Using Visual Depiction and Animation." Stanford Business School working paper. October, 1998.

Fine, Charles. *Industry Clockspeed*. 1998.

Galambos, Janos. *The Asymptotic Theory of Extreme Order Statistics*. John Wiley & Sons. 1978.

Gross, Irwin. "The Creative Aspects of Advertising." *Sloan Management Review*. Fall, 1972. pp. 83-109.

Gumbel, E. J. *Statistics of Extremes*. Columbia University Press. 1958.

Hauser, John R. "Metrics to Value R&D: An Annotated Bibliography." Working paper, International Center for Research on the Management of Technology, MIT Sloan School, March 1996.

Hauser, John R. and Florian Zettelmeyer. "Metrics to Evaluate R,D&E." *Research-Technology Management*, 40:4. July-August 1996. pp. 32-38.

Lai, T. L. and Herbert Robbins. "Maximally dependent random variables." *Proceedings of the National Academy of Science*. 73: 2, pp. 286-288. February, 1976.

Leonard-Barton, Dorothy. *Wellsprings of Knowledge*. Boston: Harvard Business School Press. 1995.

Mendelson, Haim and Ravi Pillai. "Industry Clockspeed: Measurement and Operational Implications ". *Manufacturing and Service Operations Management*, forthcoming.

Morgan, Peter and Richard Manning. "Optimal Search." *Econometrica*. Vol. 53, No. 4. July, 1985. pp. 923-944.

Pugh, Stuart. *Creating Innovative Products Using Total Design: The Living Legacy of Stuart Pugh*. Editors: Don Clausing and Ron Andrade. Addison-Wesley. 1996.

Reiss, Rolf-Dieter and Michael Thomas. *Statistical Analysis of Extreme Values*. Birkhauser Verlag. 1997.

Shoolery, Mark. Shoolery Designs. Personal conversations held from September to October, 1996.

Smith, Preston G. and Donald G. Reinertson. *Developing Products in Half the Time*. Van Nostrand Reinhold. 1995.

Srinivasan, V., William S. Lovejoy, and David Beach. "Integrated Product Design for Marketability and Manufacturing." *Journal of Marketing Research*. Volume 34. February, 1997. pp. 154-163.

Thomke, Stefan H., Eric A. von Hippel, Roland R. Franke. "Modes of Experimentation: An Innovation Process – and Competitive – Variable." Harvard Business School working paper. July, 1997.

Ulrich, Karl T. and Steven D. Eppinger. Product Design and Development. McGraw-Hill. 1995. pp. 77-104.

Vesey, Joseph T. "The New Competitors: They Think in Terms of 'Speed-to-Market'" *Academy of Management Executive*, 5:2. 1991. p. 25.

von Mises, R. "La distribution de la plus grande de *n* valuers" *Rev. Math. Union Interbalcanique*, 1. 1936. p. 141-160. Reproduced in "Selected Papers of Richard von Mises", *American Mathematical Society*, 2. 1964, pp. 271-294.

Weitzman, Martin L. "Optimal Search for the Best Alternative" *Econometrica*, 47:3. May, 1979. pp.641-654.

Wheelwright, Steven C. and Kim B. Clark. *Revolutionizing Product Development*. New York: The Free Press. 1992.