Finite scalable Bernoulli process

JR

August 1, 2016

Abstract

A special Bernoulli process is considered whose success probability scales with the number of "tokens" used each iteration. For a limited total number of tokens and a single-parameter utility function describing the scaling of marginal success utility, I derive the expected utility as a function of the number of tokens used each round.

Framework

Consider a Bernoulli process with τ iterations and probability of success $p_\tau \equiv p_1/\tau$ on each iteration. We allow the number of iterations to be in the range $\tau \in \{\lceil p_1 \rceil, \lceil p_1 \rceil + 1, \ldots, N\}$. Note that $\lceil p_1 \rceil$ is the minimum value of τ required such that $0 < p_\tau \leq 1$.

This framework approximately models the following situation: you have N "tokens" and you may choose to use t of them for each iteration of a Bernoulli process until you run out, after which the process ends. The probability of success in each Bernoulli trial is $p_t = t \times p_{t=1}$, with $t \in \{1, 2, \ldots, \min[N, \lfloor 1/p_{t=1} \rfloor]\}$. The restriction on t ensures that $0 < p_t \leq 1$. As long as $t|N$ (that is, N is an integer multiple of t), we can recover the original situation by $\tau \leftrightarrow N/t$ and $p_1 \leftrightarrow Np_{t=1}$. In the more general case, the "remainder" tokens must be dealt with separately after using t tokens for a number of trials equal to the quotient (integer part) of N/t.

The probability of having a total number of successes k in τ iterations of a Bernoulli process with success probability p_τ each time is

$$P(k, \tau) = \binom{\tau}{k} \left(\frac{p_\tau}{1 - p_\tau} \right)^k (1 - p_\tau)^\tau.$$

Suppose we assign utility u_k to getting a total number of successes equal to k. Then the expected utility from the Bernoulli process is

$$U = \sum_{k=0}^{\tau} u_k P(k, \tau).$$

For a given form of u_k, we would like to determine U as a function of τ. For instance, consider the following special cases:

- $u_k = 1$ for $k \geq k_{needed}$ and $u_k = 0$ otherwise. This applies if all you care about is getting at least k_{needed} successes. The simplest such case is $k_{needed} = 1$.
- $u_k = k$ for all k. This applies when the marginal utility of additional successes is neither increasing nor decreasing; all successes are equally valued.
- $u_k = \eta \times u_{k-1}$, with $u_0 = 0$, $u_1 = 1$. This applies when the marginal utility of each additional success is scaled by a factor of η.

We will work with the third case. Note that we can recover the first case when $k_{needed} = 1$ by letting $\eta \to 0$, and we can recover the second case with $\eta \to 1$.

For the general situation $u_k = \eta \times u_{k-1}$, we can find a closed form (rather than recursion relation) for u_k using

$$\begin{aligned} u_k &= \textstyle\sum_{i=1}^{k} \eta^{i-1} \\ &= \frac{\eta^k - 1}{\eta - 1} \qquad k \geq 1. \end{aligned} \tag{1}$$

In what follows, I simply write $p_\tau = p$, with the dependence on τ implicit. Eq. 1 implies that

$$U = \frac{\tau!(1-p)^\tau}{\eta - 1} \left[\sum_{k=1}^{\tau} \left(\frac{\eta p}{1-p} \right)^k \frac{1}{k!(\tau-k)!} - \sum_{k=1}^{\tau} \left(\frac{p}{1-p} \right)^k \frac{1}{k!(\tau-k)!} \right]. \tag{2}$$

Then we can compute sums of the required form,

$$\sum_{k=1}^{\tau} \frac{\alpha^k}{k!(\tau-k)!} = \frac{(\alpha+1)^\tau - 1}{\tau!},$$

with $\alpha \to \eta p/(1-p)$ in the first sum and $\alpha \to p/(1-p)$ in the second sum of Eq. 2. Thus

$$\begin{aligned} U &= \frac{\tau!(1-p)^\tau}{\eta-1} \left[\frac{(\eta p/(1-p)+1)^\tau - 1}{\tau!} - \frac{(p/(1-p)+1)^\tau - 1}{\tau!} \right] \\ &= \frac{\tau!(1-p)^\tau}{\eta-1} \left[\frac{(\eta p/(1-p)+1)^\tau}{\tau!} - \frac{(p/(1-p)+1)^\tau}{\tau!} \right] \\ &= \frac{(1-p)^\tau}{\eta-1} \left[(\tfrac{\eta p}{1-p} + 1)^\tau - (\tfrac{p}{1-p} + 1)^\tau \right] \\ &= \frac{(1-p)^\tau}{\eta-1} \left[\left(\frac{(\eta-1)p+1}{1+p} \right)^\tau - \left(\frac{1}{1+p} \right)^\tau \right] \\ &= \frac{1}{\eta-1} \left[((\eta-1)p+1)^\tau - 1 \right]. \end{aligned}$$

First-success criterion: $\eta = 0$

Let us analyze this first in some special cases. When $\eta \to 0$, awarding a utility of 1 for the first success and 0 for all others, we find

$$\begin{aligned} U &\to -1 \left[(1-p)^\tau - 1 \right] \\ &= 1 - (1-p)^\tau. \end{aligned}$$

The expression in the second line is exactly what we would anticipate: the expected utility is simply the probability that we get any number of successes other than 0; in other words, $U = 1 - P(0,\tau)$. Thus we want to minimize $P(0,\tau) = (1 - p_1/\tau)^\tau$. Differentiating, we find

$$\frac{dP(0,\tau)}{d\tau} = (1-p)^\tau \left(\frac{p}{1-p} + \ln(1-p)\right).$$

The first factor vanishes only when $p = 1$, while the second factor vanishes only when $p = 0$. Therefore, the slope of U with respect to τ never changes sign in the allowed domain. To check the sign of the slope, we can consider the simple case $p \to \epsilon \ll 1$, where $\frac{dP(0,\tau)}{d\tau} \to +\epsilon^2/2 > 0$. Thus $P(0,\tau)$ grows with τ and U is minimized when τ is minimized. The best-case utility is therefore

$$U_{best,\eta=0} = 1 - (1 - p_1/\lceil p_1 \rceil)^{\lceil p_1 \rceil}.$$

Note that since $U_{ideal,\eta=0} = 1$, we have $U_{best/ideal,\eta=0} = 1-(1-p_1/\lceil p_1 \rceil)^{\lceil p_1 \rceil}$. When p_1 is an integer, this ratio actually reduces to unity because in this case, each trial has a 100% chance of success.

In summary, the optimal strategy to maximize the likelihood of getting at least one success is to use as many tokens at once as possible. Given an appropriate number of tokens, the statistically optimal strategy even guarantees the ideal outcome.

Constant marginal utility: $\eta = 1$

In the case of constant marginal utility, $\eta \to 1$. In particular, let's allow approach toward $\eta = 1$ from either side by letting $\eta = 1 \pm \epsilon$ with $0 < \epsilon \ll 1$. Then the expected utility for a strategy using τ iterations is

$$\begin{array}{rcc} U & \to & \pm\frac{1}{\epsilon}\left[(1 \pm \epsilon p)^\tau - 1\right] \\ & \approx & \pm\frac{1}{\epsilon}\left[1 \pm \epsilon p \tau - 1\right] \\ & = & p\tau \\ & = & p_1, \end{array}$$

independent of the choice of τ. This is in fact what we would expect: when $\eta = 1$, the utility is just the expected number of successes, which is always $p_\tau \times \tau = p_1$. The ideal utility in this case is $U_{ideal,\eta=1} = N$, so $U_{best/ideal,\eta=1} = p_1/N$. In the "token" model, $p_1 = Np_{t=1}$ so that $U_{best/ideal,\eta=1} = p_{t=1}$ is independent of the number of tokens used per trial.

In summary, the expected utility is $U = Np_{t=1}$, independent of the number of tokens used each iteration. The statistically optimal strategy is worse than the ideal outcome by a factor equal to the individual success probability.

General case: $\eta \geq 0$

The question we are primarily interested in is: for a generic value of η, what choice of τ maximizes the expected utility U? We can solve this by maximizing the functional form of U over all (continuous) values of τ. To find candidates for local maxima, we set $dU/d\tau = 0$. We will make things a bit simpler by writing explicitly $p = p_1/\tau$ and substituting $\beta \equiv (\eta - 1)p_1$. Then

$$U = \frac{1}{\eta - 1}\left[\left(\frac{\beta}{\tau} + 1\right)^{\tau} - 1\right],$$

and

$$\begin{aligned} \frac{dU}{d\tau} &= \frac{1}{\eta-1}\frac{d}{d\tau}\left(\frac{\beta}{\tau}+1\right)^{\tau} \\ &= \frac{1}{\eta-1}\frac{\tau(1+\beta/\tau)^{\tau}}{\beta+\tau}\left((1+\beta/\tau)\ln(1+\beta/\tau) - \beta/\tau\right). \\ &= = \frac{1}{\eta-1}(1+\beta/\tau)^{\tau-1}\left((1+\beta/\tau)\ln(1+\beta/\tau) - \beta/\tau\right). \end{aligned}$$

This can only vanish when either $\tau = -\beta$ or $(1 + \beta/\tau)\ln(1 + \beta/\tau) = \beta/\tau$.

The first condition, $\tau = -\beta = (1-\eta)p_1$, is trivial: recall that $\tau \geq \lceil p_1 \rceil \geq p_1$. Thus if $\tau = (1-\eta)p_1 \geq p_1$, then $\eta \leq 0$. We have already analyzed the case when $\eta = 0$ and the conceptual significance of $\eta < 0$ is unclear, so disregard negative values of η.

The latter condition, $(1 + \beta/\tau)\ln(1 + \beta/\tau) = \beta/\tau$, occurs only when $\beta/\tau = 0$, implying either $\eta = 1$ (for which we already saw U is independent of τ) or $\tau \to \infty$. However, τ is restricted to be finite by hypothesis.

In summary, U has no critical points with respect to τ in the case $0 < \eta \neq 1$. Therefore, the optimal strategy must be either $\tau = \lceil p_1 \rceil$ or $\tau = N$. We can determine which strategy applies in each region $0 < \eta < 1$ and $1 < \eta$ separately.

To resolve these two cases can consider the simple situation in which $p = p_1/\tau \to \epsilon/|\eta - 1|$. Then $\beta/\tau \to \pm\epsilon$, with a positive sign when $\eta > 1$ and a negative sign otherwise. To leading order, one can show that $\frac{dU}{d\tau} \to \frac{1}{\eta-1}\frac{\epsilon^2}{2}$. Thus U is maximal for $\tau = N$ iff $\eta > 1$ (increasing returns) and U is maximal for $\tau = \lceil p_1 \rceil$ iff $0 < \eta < 1$ (diminishing returns).

Specifically, when $0 < \eta < 1$,

$$U_{best,0<\eta<1} = \frac{1}{\eta - 1}\left[\left((\eta - 1)\frac{p_1}{\lceil p_1 \rceil} + 1\right)^{\lceil p_1 \rceil} - 1\right].$$

The ideal utility is $\frac{\eta^N - 1}{\eta - 1}$, so

$$U_{best/ideal,0<\eta<1} = \frac{1}{\eta^N - 1}\left[\left((\eta - 1)\frac{p_1}{\lceil p_1 \rceil} + 1\right)^{\lceil p_1 \rceil} - 1\right].$$

In the special case that p_1 is an integer, $U_{best/ideal,0<\eta<1}$ reduces to $\frac{1-(\eta^N)^{p_{t=1}}}{1-\eta^N}$. This approaches unity as $p_{t=1} \to 1$ and as $\eta^N \to 0$, or equivalently as $N \to \infty$. Thus in the "token" model, the expected utility converges to the ideal utility as the number of tokens goes to infinity, given diminishing marginal utility.

When $\eta > 1$,

$$U_{best,\eta>1} = \frac{1}{\eta - 1}\left[\left((\eta - 1)\frac{p_1}{N} + 1\right)^N - 1\right]$$

and

$$U_{best/ideal,\eta>1} = \frac{1}{\eta^N - 1}\left[\left((\eta - 1)\frac{p_1}{N} + 1\right)^N - 1\right].$$

In the token model, the factor in parentheses reduces to $(\eta - 1)p_{t=1} + 1$. One can show that in this situation, $U_{best/ideal,\eta>1} \to 0$ as $N \to \infty$; the utility of the best strategy cannot "keep up" with the rapidly growing utility of the ideal outcome as more tokens are allowed.

The table below summarizes our results:

Case	Interpretation	τ_{best}	t_{best}	Best-strategy utility
$\eta = 0$	Only first success matters	Min	Max	$1 - (1 - p_1/\lceil p_1 \rceil)^{\lceil p_1 \rceil}$
$0 < \eta < 1$	Diminishing returns	Min	Max	$\frac{1}{\eta-1}\left[\left((\eta - 1)\frac{p_1}{\lceil p_1 \rceil} + 1\right)^{\lceil p_1 \rceil} - 1\right]$
$\eta = 1$	Expected number of successes	Irrelevant	Irrelevant	p_1
$\eta > 1$	Increasing returns	Max	Min	$\frac{1}{\eta-1}\left[\left((\eta - 1)\frac{p_1}{N} + 1\right)^N - 1\right]$