

Constrained Retrospective Search

Can Urgan*

Leeat Yariv^{†‡}

December 31, 2020

The search for good outcomes—be it government policies, technological breakthroughs, or a lasting purchase—takes time and effort. At times, the decision process is unconstrained: an individual seeking a well-priced product determines her search scope and time as she wishes. Often, search is constrained, either through institutions or through cognitive limitations. For instance, product-development teams often face a *design freeze*, a date at which the set of product features is locked and the first phase of research and development terminates, see [Eger, Eckert, and Clarkson \(2005\)](#). Furthermore, various grants and funding entities provide timelines that constrain the length of research. Similarly, in academia, graduate students and young faculty face research deadlines through various milestones such as early-stage paper requirements, dissertation prospectus, job-market applications, or tenure. Such limitations can also be hard-wired: going back to [Simon \(1956\)](#), the literature has often considered simple heuristics that govern individuals' search procedures.

We consider *retrospective search* in such settings: a decision maker (DM) chooses the search scope and time, selecting the best observed outcome upon stopping. We analyze the impacts of constraints when observed samples are independent and correlated.

The search literature has, by and large, focused on environments in which observed

*Department of Economics, Princeton University. Email: curgun@princeton.edu

[†]Department of Economics, Princeton University. Email: lyariv@princeton.edu

[‡]We thank Roland Benabou, Faruk Gul, Alessandro Lizzeri, Pietro Ortoleva, and Wolfgang Pesendorfer for helpful comments and suggestions. We gratefully acknowledge financial support from the National Science Foundation through grants SES-1629613 and SES-1949381.

samples are independent. Constraints in such settings have been analyzed in the celebrated work of [Stigler \(1961\)](#), who considered a DM choosing the volume of samples at the outset. We show that correlation between samples yields dramatically different implications, in terms of search scope and extent, as well as outcomes.

We consider a DM sampling from a normal distribution. The variance, or search scope, is chosen at the outset. We examine both independent samples as well as correlated samples, governed by a Brownian motion.

With independent samples, and no constraints, satisficing is optimal. When the DM selects the sample volume at the outset, the scope of search changes, and outcomes can be arbitrarily worse than those generated by search absent constraints. In contrast, when observations are correlated, optimal search entails a drawdown stopping boundary. In fact, any non-trivial satisficing generates worse outcomes than no search at all. Furthermore, the impact of constraints differs. A commitment to a search time generates a fraction of $2/\pi$ of the payoffs absent constraints, regardless of search costs.

1 Independent Samples

We start with the benchmark of independent samples. We consider discrete draws. Indeed, in continuous time, with independent draws, the DM would reach any value supported by the underlying distribution within an infinitesimal period of time.

Formally, the DM selects a search scope $\sigma \in [\underline{\sigma}, \bar{\sigma}]$ at the outset. A search scope σ is associated with a cost per sample of $c(\sigma)$, with c twice continuously differentiable, increasing, and weakly convex. At each period t in which the DM is searching, she observes an i.i.d sample $X_t^\sigma \sim N(0, \sigma)$, with density ϕ^σ and cumulative distribution function Φ^σ . We let ϕ and Φ denote the density and cumulative distribution functions of the standard normal, with standard deviation of 1, respectively. The DM has perfect recall, so keeps track of the maximum value observed, $M_t = \max_{0 \leq s \leq t} X_s^\sigma$. For simplicity, we assume that $X_0^\sigma = M_0^\sigma = 0$.

The DM's problem is then:

$$\sup_{\tau, \sigma} \mathbb{E}(M_\tau - \tau c(\sigma)). \quad (1)$$

Importantly, we assume no drift. We do so since, in most search applications, the mere passage of time does not affect discoveries' quality.

1.1 Unconstrained Search with Independent Samples

The stationarity of our environment suggests that whenever the DM continues searching with M_t , she continues searching with any $M'_t \geq M_t$. It follows that recall plays no role—the DM stops only when drawing a sufficiently high project value. The DM optimally follows a *satisficing threshold* à la [Simon \(1956\)](#). In particular, using [Robbins and Chow \(1961\)](#), one can show the following:

Proposition 1 (optimal i.i.d. stopping) *For a given search scope σ , it is optimal to stop once the satisficing threshold $S(\sigma)$ is reached, where $S(\sigma)$ solves*

$$c(\sigma) = \int_{S(\sigma)}^{\infty} (x - S(\sigma)) \phi^\sigma(x) dx.$$

Furthermore, the expected payoff $V^{iid}(\sigma)$ from using the optimal satisficing threshold $S(\sigma)$ is $V^{iid}(\sigma) = S(\sigma)$.

Intuitively, optimality of the threshold requires that it coincide with the continuation value of search. The continuation value is constant over time and, hence, $V^{iid}(\sigma) = S(\sigma)$. The characterization of the optimal threshold is then a translation of this restriction. The left hand side corresponds to the cost of an additional sample, while the right hand side corresponds to the marginal value from another draw beyond the stopping threshold.

Let $\psi(v) = \phi(v) - v \times (1 - \Phi(v))$. The function ψ can be numerically tabulated, and some analytical properties are well known: it is positive, strictly decreasing, convex, and

symmetric, with a well-defined inverse. From DeGroot (1968), the value of sampling optimally at a given search scope σ is given by

$$V^{iid}(\sigma) = \psi^{-1} \left(\frac{c(\sigma)}{\sigma} \right) \sigma.$$

The optimal search scope can be readily calculated as the maximizer:

Corollary 1 (optimal i.i.d. search scope) *The optimal search scope σ maximizes $\psi^{-1} \left(\frac{c(\sigma)}{\sigma} \right) \sigma$.*

As it turns out, even if the DM could select a search scope freely at any period, the constant search identified in Corollary 1 would be optimal. Intuitively, as recall plays no role, the DM faces an identical optimization problem in each period while she searches.

When costs are sufficiently convex, say log-convex, the optimal search scope is extremal, either $\underline{\sigma}$ or $\bar{\sigma}$.

1.2 Pre-Committed Time with Independent Samples

Consider now the case in which the DM simultaneously commits to the search scope and the number of samples she draws. For a fixed search scope, this case resembles that studied by Stigler (1961).¹

Suppose the DM selects a search scope of σ and n samples of values. The resulting payoff would then be the highest order statistic from a sample of n normal variables censored at 0, net of the overall costs:

$$\sigma \int_{-\infty}^{\infty} \max(x, 0) n \phi(x) \Phi^{n-1}(x) dx - c(\sigma)n$$

The integral term is just the expected maximum of n draws from a censored normal distribution with $\sigma = 1$.² We denote the random variable corresponding to the highest

¹Stigler (1961) did not consider an outside option for the DM, essentially assuming she selects a non-trivial sample and potentially accepts arbitrarily bad outcomes. Our restriction that $X_0 = M_0 = 0$ therefore changes somewhat the calculus underlying the characterization of the optimal policy, if not its description.

²The distribution of a normal variable censored at 0 is still scale-invariant, see the Online Appendix.

order statistic of n normal samples censored at 0 with $\sigma = 1$ by $Y_{(n)}$. Its expectation can be tabulated numerically, but analytical formulations are challenging for $n > 3$. Nonetheless, it can be shown that it is increasing and concave in n .

Proposition 2 (optimal i.i.d. constrained search) *For any search scope σ , the optimal number of samples n solves:*

$$\mathbb{E}(Y_{(n+1)}) - \mathbb{E}(Y_{(n)}) < \frac{c(\sigma)}{\sigma} < \mathbb{E}(Y_{(n)}) - \mathbb{E}(Y_{(n-1)}).$$

For any sample number n , the optimal search scope, if interior, solves $\mathbb{E}(Y_{(n)}) = c'(\sigma)n$.

The first restriction corresponds to the analysis of [Stigler \(1961\)](#). The DM selects the maximal number of samples n such that the marginal benefit of the n 'th sample is exceeds its cost, while the marginal benefit of the $(n + 1)$ 'th sample does not.

The second restriction corresponds to a first-order condition with respect to the search scope. Importantly, the optimal search scope in the optimal dynamic search, characterized in [Corollary 1](#) differs from that selected by a DM constrained to a pre-committed number of samples.

2 Correlated Samples

In many applications, innovation begets innovation. In research and development, one new technique or idea builds on previous ones. In geological surveys, one plot's mineral returns are indicative of the returns from an adjacent plot. In online shopping, suggested items by commerce platforms are often associated with prior considered items. As we show, the correlation of samples yields very different conclusions than those gleaned from the i.i.d. case.

We model correlation over time using a Brownian motion governing the path of values, similar to [Callander \(2011\)](#). For any search scope σ , the DM observes at time t

the value X_t^σ satisfying

$$dX_t^\sigma = \sigma dB_t,$$

where B_t is the standard Brownian motion with standard deviation of 1. As before, the DM selects the search scope $\sigma \in [\underline{\sigma}, \bar{\sigma}]$ at the outset, with the cost $c(\sigma)$ defined as before. The DM has perfect recall so records the maximum value observed at any time t , namely $M_t = \max_{0 \leq s \leq t} X_s^\sigma$. We continue assuming that $X_0^\sigma = M_0^\sigma = 0$. The DM then faces the same problem as specified in (1), namely $\sup_{\tau, \sigma} \mathbb{E}(M_\tau - \tau c(\sigma))$.

We continue assuming no drift both to match most search applications and to maintain comparability with the independent-sample case. We note, however, that in this setting, the maximum value itself exhibits a form of drift. Namely, with search scope σ , at any time t , we have $\mathbb{E}(M_t) = \sigma \sqrt{2t/\pi}$.

2.1 Unconstrained Search with Correlated Samples

The analysis of the optimal search policy when observations are correlated follows that of [Urgun and Yariv \(2020\)](#). Perfect recall is now important. With any observed value, the DM assesses the time it would take to reach a value exceeding the maximum value observed, which she can collect immediately. With correlated samples, a current low observed value *relative* to the historical maximum suggests a long time, entailing high search costs, for search to pay off. In particular, the optimal stopping policy is now identified by a *stopping boundary*, which may depend on the recorded maximum.

The following proposition characterizes the optimal stopping boundary, which turns out to be a *drawdown stopping boundary*. That is, the DM stops searching whenever the observed value is a fixed distance below the recorded maximum. That fixed distance is referred to as the *drawdown size*.

Proposition 3 (Urgun and Yariv, 2020) *For any search scope σ , the DM stops searching as*

soon as $X_t^\sigma \leq M_t^\sigma - d^\sigma$ at any time t , where the drawdown size d^σ is identified by $d^\sigma = \frac{\sigma^2}{2c(\sigma)}$. The expected payoff is given by $V^{corr}(\sigma) = \frac{d^\sigma}{2} = \frac{\sigma^2}{4c(\sigma)}$.

The optimal search scope then maximizes $\sigma^2/c(\sigma)$, so that we have:

Corollary 3 (optimal correlated search scope) *The optimal search scope, if interior, solves*

$$\frac{2c(\sigma)}{c'(\sigma)} = \sigma.$$

In fact, [Urgun and Yariv \(2020\)](#) show that such a constant search scope is optimal even when scope can be adjusted dynamically. When costs are log-convex, an interior solution is unique and exhibits natural comparative statics: as costs become more log-convex, the optimal search scope declines. Importantly, the optimal search scope in the correlated case differs from the optimal search scope in the independent case.

2.2 Pre-Committed Time with Correlated Samples

Suppose search is to take place over a period of time T , analogous to the pre-committed number of samples considered in the independent-sample case. If the DM uses a search scope σ , the expected payoff is:

$$\bar{V}^{corr}(\sigma; T) = \mathbb{E}(M_T) - c(\sigma)T.$$

Following standard arguments, the record-high level M_t at time t has the same distribution as that of $|X_t|$. The time- T realization of B_T is normally distributed with mean 0 and variance T . Therefore, $\mathbb{E}(|B_t|) = \sqrt{2T/\pi}$. Since $X_t = \sigma B_t$, we have that:

$$\bar{V}^{corr}(\sigma; T) = \sigma\sqrt{2T/\pi} - c(\sigma)T.$$

In particular, we can solve for the optimal search horizon and scope:

Proposition 4 (optimal correlated constrained search) *For any search scope σ , the optimal*

fixed search time is $T^\sigma = \frac{\sigma^2}{2\pi(c(\sigma))^2}$. The resulting expected payoff is $\bar{V}^{corr}(\sigma; T^\sigma) = \frac{\sigma^2}{2\pi c(\sigma)}$.

The optimal search scope, if interior, solves $\frac{2c(\sigma)}{c'(\sigma)} = \sigma$.

Unlike the independent-sample case, the optimal search scope is the same, whether or not the DM is constrained in her search.

2.3 Satisficing with Correlated Samples

When samples are independent, we showed that an unconstrained DM uses a satisficing threshold to govern her stopping policy. How would satisficing perform when samples are correlated? We now show that, when observations are correlated, the DM would prefer to stop immediately rather than search following *any* positive satisficing threshold.

Indeed, suppose the DM stops whenever she reaches a satisficing level $S \geq 0$. That is, the DM stops only when hitting a new maximum of S , which is then her payoff. Let $\tau_S^{sat} = \inf\{t \geq 0 : X_t = S\}$. That is, τ_S^{sat} is the random time at which a satisficing retrospective searcher stops. Since the underlying process generating outcomes X_t has no drift, we have that for any $S > X_0$, $P(\tau_S^{sat} < \infty) = 1$, but $\mathbb{E}(\tau_S^{sat}) = \infty$. Namely, the expected time it takes a driftless Brownian motion to hit a certain threshold above its starting point is infinite. The DM's expected payoff is then negative for any bounded $S > 0$. In particular, the optimal satisficing level is $S = X_0 = 0$, indicating that a satisficing DM stops immediately.³

Why is satisficing so inefficient for a retrospective searcher? While the DM may get a high level of utility upon stopping, determined by her satisficing level, she may also continue her search when low values are observed, ones that would ideally induce her to cease search. That possibility is prohibitively costly in expectation. A natural extension to the satisficing heuristic would then allow for a threshold *below* which the DM stops.⁴

³Similar extreme predictions occur with drift. Indeed, suppose values exhibit a drift of μ and that the search scope σ comes at a flow cost of c . The expected payoff from a satisficing threshold S would be $S - Sc/\mu$. Thus, for low costs such that $c < \mu$, the optimal satisficing level would be infinite, while for high costs such that $c > \mu$, optimal satisficing would dictate immediate stopping.

⁴With scope σ , expected payoffs from satisficing threshold $S \geq 0$ and departure threshold $D \leq 0$ are:

$$\Pi^{D,S}(\sigma) = \frac{|D|S}{|D|+S} \left(\frac{|D|}{|D|+S} + \ln \left(\frac{|D|}{|D|+S} \right) \right) - \frac{|D|S}{\sigma^2} c(\sigma).$$

One can show that, optimally, $|D| = S$.

3 Impacts of Constraints

Search with independent samples generates different behavior than search with correlated samples. While satisficing is optimal with independent samples, it does poorly with correlated samples, when a drawdown stopping boundary is optimal. The optimal search scope also responds differently to features of the cost. With independent samples, unconstrained search leads to the use of the *last* value observed, while constrained search leads to the last value being used with *some probability*, which vanishes as the number of samples grows. In contrast, with correlated samples, the last value observed is *never* used.⁵

The impacts of constraints also differ across the two settings. With independent samples, the optimal search scope changes when search time is chosen at the outset; with correlated samples, the optimal search scope is the same with and without constraints.

Certainly, in both settings, constraints reduce the value of search. When samples are correlated, committing to a search time entails the loss of a fixed fraction of the search value. In contrast, as we now show, with independent samples, ex-ante commitment to a search time may have severe consequences, depending on search costs.

For simplicity, consider a linear cost function: $c(\sigma) = a\sigma$ with $a > 0$. In this case, it is easy to verify that, in both settings, with or without constraints, the optimal search scope is $\bar{\sigma}$. Denote the corresponding expected payoff from the optimal policy in the independent setting for the unconstrained and constrained search by $V^{iid}(a)$ and $\bar{V}^{iid}(a)$, and for the correlated setting by $V^{corr}(a)$ and $\bar{V}^{corr}(a)$, respectively. An application of our results thus far is then the following:

Corollary 4 (constraints and costs) *For the independent-sample setting, $\lim_{a \rightarrow 0} V^{iid}(a) / \bar{V}^{iid}(a) = \infty$. For any $a > 0$, in the correlated-sample setting, $V^{corr}(a) / \bar{V}^{corr}(a) = \frac{2}{\pi}$.*

⁵For the constrained case, this is an artifact of our continuous-time setting: even when search is constrained to horizon $T > 0$, the event that the maximal value occurs at T has zero probability.

Taken together, our results illustrate the dramatically different behaviors independence and correlation across samples generate. In particular, the impacts of constraints can differ drastically across the two environments.

References

- CALLANDER, S. (2011): "Searching and learning by trial and error," *American Economic Review*, 101(6), 2277–2308.
- DEGROOT, M. H. (1968): "Some problems of optimal stopping," *Journal of the Royal Statistical Society: Series B (Methodological)*, 30(1), 108–122.
- EGER, T., C. ECKERT, AND P. J. CLARKSON (2005): "The role of design freeze in product development," in *DS 35: Proceedings ICED 05, the 15th International Conference on Engineering Design, Melbourne, Australia, 15.-18.08. 2005*.
- ROBBINS, H., AND Y. CHOW (1961): "A Martingale System Theorem and Applications," in *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability*. University of California Press.
- SIMON, H. A. (1956): "Rational choice and the structure of the environment.," *Psychological review*, 63(2), 129.
- STIGLER, G. J. (1961): "The economics of information," *Journal of political economy*, 69(3), 213–225.
- URGUN, C., AND L. YARIV (2020): "Retrospective Search: Exploration and Ambition on Uncharted Terrain," *mimeo*.