SIMULTANEOUS SEARCH

BY HECTOR CHADE AND LONES SMITH[1]

We introduce and solve a new class of "downward-recursive" static portfolio choice problems. An individual simultaneously chooses among ranked stochastic options, and each choice is costly. In the motivational application, just one may be exercised from those that succeed. This often emerges in practice, such as when a student applies to many colleges or when a firm simultaneously tries several technologies.

We show that such portfolio choice problems quite generally entail maximizing a submodular function of finite sets—which is NP-hard in general. Still, we show that a greedy algorithm finds the optimal set, finding first the best singleton, then the best single addition to it, and so on. We show that the optimal choices are "less aggressive" than the sequentially optimal ones, but "more aggressive" than the best singletons. Also, the optimal set in general contains gaps. We provide some comparative statics results on the chosen set.

KEYWORDS: Search, portfolio choice, submodular, greedy algorithm.

1. INTRODUCTION

WE INTRODUCE AND SOLVE a new class of "downward-recursive" portfolio choice problems. For instance, a decision maker (DM) simultaneously chooses among ranked stochastic options, each choice is costly, and only the best realized alternative is exercised.

Our paper generalizes Stigler (1961), who analyzes optimal static wage search. Unlike Stigler, we do not assume a priori identical prizes, and we characterize both the optimal sample size and *choice composition*. Weitzman (1979) also explores a problem with a priori distinct prizes, but in the sequential world. His is a nice application of Gittins' solution of the bandit problem. Each option can be assigned an index in isolation of all others; sequentially, one simply chooses the unexplored option with the highest index.

In our problem, no such simple index rule presents itself. Instead, we find ourselves faced with the maximization of a submodular function of sets of alternatives—to be sure, a complex combinatorial optimization problem. Nevertheless, we show that an economically natural algorithm produces the optimal set in a quadratic number of steps.

We then explore properties of the optimal set. We ask, for instance, how much risk should one take. We show that the optimal portfolio is more aggressive than the set of best options taken individually, but less aggressive than Weitzman's sequential choices.

We also ask how varied should the choices be. We argue in favor of an upwardly diverse portfolio: For a rich enough array of possible options and low enough costs, a connected "interval" of similarly risky prospects is not optimal. We next provide a key comparative static, showing how the choice set improves when acceptance chances rise and the acceptance chances of better alternatives rise proportionately more.

We believe that our problem is not without substantive practical value.

EXAMPLE 1: A student must make a costly and simultaneously application to several colleges, and is accepted with smaller chances by the better schools.

EXAMPLE 2: A large firm wishes to choose a technology: several are available and all are costly to explore; some will work out and others will not. Finally, the firm is in a hurry (e.g., it is in a race with other firms) and must simultaneously choose which to explore.

EXAMPLE 3: An economics department must pay travel expenses for new Ph.D. job candidates. Each school ranks the candidates; the better ones are harder to hire.

Our paper may also be more topically viewed as a foundation for the recent literature on directed search (e.g., allowing employees to choose which jobs to apply to). We solve this decision problem for multiple applications and heterogeneous jobs.[2]

We first describe the problem. We introduce the algorithm and prove its optimality. We then explore the properties of the optimal set: Does the DM insure herself or gamble? Are the optimal choices similar or disparate? What if success rates increase? The Appendix contains one more algebraic proof. We also refer the reader to the working paper version (Chade and Smith (2005)) for additional results, examples, and discussion.

2. THE PORTFOLIO PROBLEM

A decision maker (DM) can consume prizes from a finite set $N = \{1, 2, \ldots, N\}$. Here, N is a natural number, but abusing notation, we denote this set also by N and denote its subsets by 2^N (with the subset inclusion order). Let $f: 2^N \mapsto \mathbb{R}_+$ be a strictly increasing function with $f(\emptyset) = 0$. Interpret $f(S)$

---

[2]See Burdett, Shi, and Wright (2001) and Albrecht, Gautier, and Vroman (2002). Perhaps the first equilibrium paper with multiple simultaneous searches is Burdett and Judd (1983).

as the expected value of subset S and put $z_i \equiv f(i) > 0$.[3] Prizes are random and the prize set S has failure chance $\rho(S) \in [0, 1)$ for all $S \neq \emptyset$ (and $\rho(\emptyset) = 1$). Because $\alpha_i \equiv 1 - \rho(i)$ is the success chance of prize i, the ex post payoff is $u_i \equiv f(i)/\alpha_i$. We assume that prize 1 is ex post the best, prize 2 is the next, and so forth so that $u_1 > u_2 > \cdots > u_N$.

Say that U is *above* L, written $U \sqsupseteq L$, if the worst prize in U beats the best in L. We assume that the portfolio S is worth less than the sum of its parts. Specifically, this payoff function is *downward recursive* (DR), so that for all sets $U \sqsupseteq L$ in N,

$$(1) \qquad f(U + L) = f(U) + \rho(U)f(L).$$

We observe that ρ is multiplicative in a DR payoff structure, because for all $U \sqsupseteq M \sqsupseteq L$,

$$f(U + M + L) = f(U + M) + \rho(U + M)f(L)$$
$$= [f(U) + \rho(U)f(M)] + \rho(U)\rho(M)f(L)$$

so that $\rho(U + M) = \rho(U)\rho(M)$. Because $\rho < 1$ and is multiplicative, ρ *is strictly decreasing*.

The cost of a portfolio S is given by a function $c(|S|)$, where S has cardinality $|S|$, $c(0) = 0$, and c is increasing and convex on the nonnegative integers. We assume $z_i > c(1)$ for all i, thereby pruning weakly dominated prizes. This paper studies a one-shot maximization of $v(S) = f(S) - c(|S|)$. Of course, $v(\emptyset) = 0$. Our analysis will frequently require consideration of finite subdomains $D \subseteq N$. For such $D \subseteq N$, let $\Sigma^*(D)$ solve

$$(2) \qquad \max_{S \subseteq D} v(S)$$

and denote $\Sigma^* \equiv \Sigma^*(N)$.

We also explore two prominent special cases of (2). By a *fixed cost per application*, we mean that $c(|S|) = \bar{c}|S|$ for some $\bar{c} > 0$. In the *fixed sample size* n case, $c(|S|) = 0$ if $|S| \leq n$ and $c(|S|) = \infty$ if $|S| > n$. Problem (2) becomes $\max_{S \subseteq D \& |S| = n} f(S)$ with solution $\Sigma_n(D)$. We then define $\Sigma_n \equiv \Sigma_n(N)$. Notice that $\Sigma^* = \Sigma_n$ for some n.

3. APPLICATIONS

A. The College Problem

A student must choose once and for all a portfolio $S \subseteq N$ of colleges to which to apply for admission, at cost $c(|S|)$. The best is 1, the second best is 2, and

---

[3] We avoid set notation, writing $i = \{i\}$, $A + B = A \cup B$, $A - B = A \setminus B$, $(i, j) = \{k \in N | i < k < j\}$, etc.

so on. The student's cardinal utility (*ex post payoff*) from attending college i is u_i, where $u_1 > u_2 > \cdots > u_N$. Her *admission chance* at college i is $\alpha_i \in (0, 1]$. Intuitively one might imagine the *inverse ordering* $\alpha_1 < \alpha_2 < \cdots < \alpha_N$, but this is inessential, as we shall see. The acceptance decisions by any set of colleges are independent. For instance, this arises when colleges perceive noisy conditionally independent and identically distributed signals of a student's caliber and she fully knows her true caliber. The expected payoff of college i *alone* is $z_i = \alpha_i u_i$.

Working recursively, either one gets into the best college in S or one does not; if rejected, one either gets into the next best or not, etc. Because $\rho(S) \equiv \prod_{i \in S}(1 - \alpha_i)$ is the chance of rejection by all colleges in the set S, the *gross payoff* may be decomposed as

$$(3) \qquad f(S) = \sum_{i=1}^{|S|} z_{(i)} \prod_{\ell=1}^{i-1}(1 - \alpha_{(\ell)}) = \sum_{i=1}^{|S|} z_{(i)} \rho_{(i-1)}(S),$$

where (i) is the ith best-ranked college in the set S, so that $z_{(i)} \equiv \alpha_{(i)} u_{(i)}$, and $\rho_{(i-1)}(S) = \prod_{\ell=1}^{i-1}(1 - \alpha_{(\ell)})$ is the chance of being rejected by the top-ranked $i-1$ schools in set S.

This college structure contains the generality of the DR payoff structure of Section 2 and, for definiteness, we sometimes cast our results in the language of this application.

B. Other Singleton Prize Models

The technology choice clearly has this structure. Hiring at the economics department assumes this form after some reworking. Indeed, assume that (i) interviews are costly, but do not inform the hiring decision; (ii) each department needs at most one job candidate; (iii) after the interview stage, the market clears top to bottom, so that the better recruits are available with smaller chance to any school below the top.

C. Correlated Rejection Chances

Modify the college problem so that rejection from school i scales down the acceptance chance at colleges $j > i$ by a factor $\beta_i \in [0, 1]$. Portfolio S then has value $f(S) = \sum_{i=1}^{|S|} z_{(i)} \prod_{l=1}^{i-1}(1 - \alpha_{(l)})\beta_{(l)}$. This derives from a consistent probability distribution over N for large enough β_i (all i) and reduces to (2) when $\beta_i \equiv 1$. Because of the DR structure, an equivalent college problem exists with independent admission events: Assume acceptance rates $\bar{\alpha}_i = 1 - (1 - \alpha_i)\beta_i$ and college payoffs $\bar{u}_i = \alpha_i u_i / \bar{\alpha}_i$. The ex post payoffs $\bar{u}_i$ fall in i also when the β_i's are large enough.

D. *One-Shot Multidecisions for Dynamic Choice with Payoff Discounting*

The DM enjoys payoffs from all successful options, but can only consume one per period. He thus "eats" the best first, etc. Future payoffs are discounted by the factor $\delta \in [0, 1)$. One can show that the expected payoff of portfolio S is $f(S) = \sum_{i=1}^{|S|} z_{(i)} \prod_{l=1}^{i-1}(1 - \alpha_{(l)} + \alpha_{(l)}\delta)$. Here, an equivalent college problem requires $\bar{\alpha}_i = (1 - \delta)\alpha_i$ and $\bar{u}_i = u_i/(1 - \delta)$.

4. THE SOLUTION

4.1. *Consistency Checks on the Optimal Set*

Computing the optimal set is a complex task, but we are able now to provide two useful tests that it must obey. The DR equality (1) implies a key *ordinal* property—*downward maximization*, that is, optimizations on sets imply optimizations on lower ends of those sets:

LEMMA 1: *Let* $\Sigma_n = U + L$, *where* $U \sqsupseteq L$ *and* L *has* k *elements. Then* $\Sigma_k(D) = L$, *where* D *are those options in* N *not better ranked than the best in* L.

Let $MB_k(S) \equiv f(S + k) - f(S)$ be the marginal benefit of adding k to $S \subset N - k$. If ex ante and ex post ranks of options agree, their marginal values are likewise ranked.

LEMMA 2: *Assume* $z_i > z_j$ *and* $i < j$. *Then the marginal benefits of* i, j *are ordered* $MB_i(S) \equiv f(S + i) - f(S) > f(S + j) - f(S) = MB_j(S)$ *for any set* $S \subset N \setminus \{i, j\}$.

PROOF: Since $i < j$, we may write $S = U + M + L$ for sets (upper) $U = [1, i) \cap S$, (middle) $M = (i, j) \cap S$, and (lower) $L = (j, N] \cap S$. So $U \sqsupseteq M \sqsupseteq L$. Consider the *suboptimal* implementation policy for $S + i$: Accept the best available option unless it is i, in which case accept the best option in M (if available) over i. So by (1),

$$f(S + i) \geq f(U) + \rho(U)\big(f(M) + \rho(M)[z_i + (1 - \alpha_i)f(L)]\big)$$
$$> f(U) + \rho(U)\big(f(M) + \rho(M)[z_j + (1 - \alpha_j)f(L)]\big)$$
$$= f(S + j),$$

because $z_i - \alpha_i f(L) > z_j - \alpha_j f(L)$, given $z_i = \alpha_i u_i > \alpha_j u_j = z_j$ and $u_i > u_j > f(L)$. *Q.E.D.*

If $j \in \Sigma_n(N)$, then setting $S = \Sigma_n(N) - j$ and using the inequality $f(S + i) > f(S + j)$ yields at once a simple insight into Σ^*. For any chosen option, any better ranked one with greater expected payoff is also chosen:

LEMMA 3: *Assume* $z_i > z_j$ *and* $i < j$. *If* $j \in \Sigma_n(N)$, *then* $i \in \Sigma_n(N)$.

4.2. *An Optimal Marginal Improvement Algorithm*

A greedy algorithm at each step makes the locally optimal choice in the hope of finding the global optimum. The following greedy algorithm identifies Σ^* via an inductive procedure:

MARGINAL IMPROVEMENT ALGORITHM (MIA): Let $Y_0 = \emptyset$.

Step 1: *Choose any* $i_n \in \arg\max_{i \in N \setminus Y_{n-1}} f(Y_{n-1} + i)$.

Step 2: *If* $f(Y_{n-1} + i_n) - f(Y_{n-1}) < c(n) - c(n-1)$, *then stop*.

Step 3: *Set* $Y_n = Y_{n-1} + i_n$ *and go to Step* 1.

So one first identifies the option i_1 whose expected payoff z_i is largest.[4] At any stage n, one finds the option i_n that affords the largest marginal benefit over the college set constructed so far. The algorithm stops if the *net* marginal benefit turns negative.

THEOREM 1: *The MIA identifies the optimal set Σ^* for problem* (2) *with $D = N$.*

Standard combinatorial optimization proofs proceed by policy improvement. Our proof by a double induction below on n and N is mostly different, exploiting two properties of DR functions. We have the ordinal property of $\Sigma_n(D)$ in Lemma 1, and the following cardinal property.

Let $Y_n(D)$ be the n options chosen by the MIA from *domain* $D \subset N$. In set D, let us suggestively call the best option 0. Let options $i, j \in N - D$ be lower ranked than 0. Then

$$(4) \qquad f(D + i) - f(D + j) = \rho(0)[f(D - 0 + i) - f(D - 0 + j)].$$

Thus, if the marginal benefit of i exceeds j given $D - 0$, then it does so given D. So:

LEMMA 4: *We have* $|Y_n(D - 0) \cap Y_n(D)| \geq n - 1$ *and* $|Y_{n-1}(D - 0) \cap Y_n(D)| = n - 1$.

---

[4]The proof actually ignores the nongeneric possibility of tied values of multiple arg max. With tied values, there exists a vanishing sequence of ε payoff perturbations that renders Σ_n uniquely optimal along the ε sequence. By the Theorem of the Maximum, this constant solution correspondence of the ε-perturbed problems gives the solution of the unperturbed limit problem. So the choice Σ_n is optimal.

PROOF: To see both claims, observe that if $0 \notin Y_n(D)$, then $Y_n(D) = Y_n(D - 0) \supset Y_{n-1}(D - 0)$, but if $0 \in Y_n(D)$, then $Y_n(D) = 0 + Y_{n-1}(D - 0)$ by (4).   Q.E.D.

To show $\Sigma_n(D) = Y_n(D)$, we find (a) downward maximization (Lemma 1) links $\Sigma_n(D)$ and $\Sigma_{n-1}(D)$, as well as $\Sigma_n(D)$ and $\Sigma_n(D - 0)$; (b) the marginal benefit property (Lemma 4) and the MIA likewise link $Y_n(D)$ as well as $Y_{n-1}(D)$, and $Y_n(D)$ and $Y_n(D - 0)$; and (c) $\Sigma_{n-1}(D) = Y_{n-1}(D)$ and $\Sigma_n(D - 0) = Y_n(D - 0)$, by induction assumption.

PROOF OF THEOREM 1: We first show that $Y_n(D) = \Sigma_n(D)$ if $D \subseteq N$ and $n \leq |D|$. This is trivial for $|D| = 1$. Assume it for $\hat{D} \leq D$, $\hat{n} \leq n$, not both with equality. If $i \notin Y_n(D) \cup \Sigma_n(D)$, then $\Sigma_n(D) = \Sigma_n(D - i) = Y_n(D - i) = Y_n(D)$ by induction. So let $Y_n(D) \cup \Sigma_n(D) = D$.

CASE 1—$0 \in \Sigma_n(D)$: By Lemma 1, $\Sigma_n(D) = 0 + \Sigma_{n-1}(D - 0)$, so that 0 yields the greatest marginal improvement to $\Sigma_{n-1}(D - 0) = Y_{n-1}(D - 0)$ by the induction assumption. If $0 \notin Y_{n-1}(D)$, then $Y_{n-1}(D - 0) = Y_{n-1}(D)$ and hence $Y_n(D) = 0 + Y_{n-1}(D) = \Sigma_n(D)$. If $0 \in Y_{n-1}(D) \subset Y_n(D)$, then Lemma 4 yields $Y_n(D) = 0 + Y_{n-1}(D - 0) = \Sigma_n(D)$.

CASE 2—$0 \notin \Sigma_n(D)$: By induction and the case premise, $Y_n(D - 0) = \Sigma_n(D - 0) = \Sigma_n(D)$. Also, $Y_n(D)$ and $Y_n(D - 0)$ overlap on $n - 1$ or n schools by Lemma 4. Because $Y_n(D) \cup \Sigma_n(D) = D$, we have $\Sigma_n(D) = D - 0$ and $Y_n(D) = D - k$ for some k. Assume $k \neq 0$ or we are done. Thus, $\Sigma_{n-1}(D) = Y_{n-1}(D) = D - j - k$ for some $j \neq k$ by the induction assumption and the MIA. If $j = 0$, then the best addition to $Y_{n-1}(D) = D - 0 - k$ is 0 and so $\Sigma_n(D) = D - k$. Contradiction. Assume $j \neq 0$ and without loss of generality let $k < j$. Write $D = U + L$, where $U = D \cap [1, k]$ and $L = D - U$, so $U \sqsupseteq L$. If $\alpha_0 > \alpha_j$, then easily $z_0 > z_j$ and so $f(D - j) > f(D - 0)$ by Lemma 2. Hence, $\Sigma_n(D) = D - 0$ is suboptimal, a contradiction. In addition, when $\alpha_0 < \alpha_j$,

$$(5) \qquad f(D - j) = f(U - j - k) + \rho(U - j - k)[f(k) + \rho(k)f(L)]$$
$$= f(D - j - k) + \rho(U - j - k)\big[f(k) - (1 - \rho(k))f(L)\big]$$

exceeds $f(D - 0 - k) + \rho(U - 0 - k)[f(k) - (1 - \rho(k))f(L)] = f(D - 0)$, given $f(D - j - k) > f(D - 0 - k)$, $\rho(U - j - k) > \rho(U - 0 - k)$, and $f(k)/(1 - \rho(k)) = u_k > f(L)$.

Because the cost of a portfolio depends *only* on its size, $\Sigma^*(N) = \Sigma_n(N)$ for some n. The stopping rule is optimal because the cost $c(n)$ is convex in n and because f has *diminishing returns*—$f(S + k) - f(S)$ is decreasing in S for any $k \notin S \subseteq N$—as we see below.   Q.E.D.

To see diminishing returns, let us write the marginal benefit of adding college k to a set S as follows. Partition $S = U_k + L_k$, where $U_k = [1, k) \cap S$ and $L_k = (k, N] \cap S$. Then (5) yields

$$(6) \qquad MB_k(S) = f(S + k) - f(S) = \rho(U_k)[z_k - \alpha_k f(L_k)].$$

Because $u_k > f(L_k)$, and ρ is decreasing and f increasing, we have the following result:

LEMMA 5: *Any DR function $f : 2^N \mapsto \mathbb{R}$ has diminishing returns.*

Intuitively, additions to the current portfolio grow less valuable as more options are added. Note that, by Lemma 5, $v(S + i) - v(S) < v(i) - v(\emptyset) = z_i - c(1)$ whenever $S \neq \emptyset$. So choosing all options with $z_i > c(1)$ yields a suboptimally large portfolio.

4.3. *Submodular Optimization*

As noted, the value of a portfolio is less than the sum of its parts, because each option exerts a negative externality on the others. To cleanly capture this notion, call a function f on 2^N submodular if $f(S \cap T) + f(S \cup T) \leq f(S) + f(T)$ for any two subsets S and T of N.

LEMMA 6: *Any DR function f is submodular and thus so is $v : 2^N \mapsto \mathbb{R}$ in (2).*

PROOF: First, f is submodular because it has diminishing returns (Proposition 1.1 in Lovász (1982)).[5] Next, $-c(|S|)$ is a concave function and, therefore, is a submodular function (Proposition 5.1 in Lovász (1982)).[6] *Q.E.D.*

It is well known that the maximization of a *general* submodular set function is NP-hard and thus computationally intractable. Indeed, no polynomial algorithm exists for it—this is independent of the P $\neq$ NP problem; see Lovász (1982, p. 252). By exploiting the special functional form of our objective function v, the MIA quickly finds the optimal set Σ^* for all DR submodular functions. One must, in principle, calculate the values of all 2^N college application patterns. Yet our algorithm succeeds in polynomial time: Initially, one examines N options and finds the best one. One then examines the next $N - 1$ and finds the second best, etc. This amounts to $\sum_{i=0}^{N-1}(N - i) = N(N + 1)/2 = O(N^2)$ steps.

Let us step back and ask whether the MIA's success was predestined in light of the recent theory of combinatorial optimization. One can show that $f(S)$

---

[5]Gul and Stacchetti (1999) recently used this property in the economics literature. See related work by Kelso and Crawford (1982) on the gross substitutes condition.
[6]Observe that a sum of submodular functions, like $f + (-c)$, is submodular.

is (what is known as) *semistrictly quasiconcave*.[7] As with standard quasiconcavity, local then implies global optimization. It does not, however, imply that a "steepest ascent" algorithm like the MIA will succeed, as we prove it does for the class of DR payoff functions.[8]

5. PROPERTIES OF THE OPTIMAL SET

5.1. *Aggressiveness of the Optimal Choices*

How "risk-taking" should the portfolio be? To flesh this out, we employ *vector* first order stochastic dominance (FSD). The set $S \subseteq N$ is *more aggressive* than the same-size set $S' \subseteq N$ in the sense of FSD when $s_{(i)} \leq s'_{(i)}$ for all i, where $s_{(i)}$ is the ith best school in S and $s'_{(i)}$ in S'. Write this as $S \succeq S'$ and as $S \succ S'$ if also $S \neq S'$. Thus, $\{1, 2\} \succ \{2, 3\}$.

We now compare the best set Σ^* against two easily computed benchmarks.

A. Portfolio choices are more aggressive than top singletons

Consider the set $Z_{|\Sigma^*|} \subseteq N$ of options with the $|\Sigma^*|$ highest expected payoffs $z_i = \alpha_i u_i$. Unlike the portfolio Σ^*, this set ignores the web of cross-college external effects, as captured by (3).

THEOREM 2: *The best portfolio Σ^* is more aggressive than the best singletons $Z_{|\Sigma^*|}$.*

PROOF: It suffices to show that if $i < j$ and $z_i > z_j$, then the MIA picks i before j. By Lemma 2, for *any* portfolio S excluding i, j, we have $MB_i(S) > MB_j(S)$.   *Q.E.D.*

For an intuition, consider expression (3) for expected payoffs, written as $\sum_i z_{(i)} \rho_{i-1}(S)$. If options in Σ^* do not have the highest z_i's, then they must compensate with a higher $\rho_{i-1}(S)$. Thus acceptance chances are lower and these options must be better ranked.

To see that the order can be strict, assume three colleges with $\alpha_1 = 0.1$, $\alpha_2 = 0.9$, $\alpha_3 = 1$, $u_1 = 1$, $u_2 = 0.5$, and $u_3 = 0.48$. Notice that $z_3 = 0.48 > z_2 = 0.45 > z_1 = 0.1$. One can show that $\Sigma_2(3) = \{1, 3\}$, which is strictly more aggressive than $Z_2 = \{2, 3\}$.

Static portfolio maximization thus precludes "safety schools." One never applies to a school for its high admissions rate, when not otherwise justified by its expected payoff. However, one might apply to a high-ranked "stretch school," despite the low expected payoff.

---

[7]See Murota and Shioura (2003). Precisely, $f(S)$ only satisfies their property ($SSQM_w$) on page 472. See Chade and Smith (2005).

[8]In Chade and Smith (2005), we show that the MIA works for some non-DR functions too.

The "no safety school" substance of Theorem 2 is undermined with arbitrary correlation (not of the form in Section 3C). For an extreme example with perfect correlation, assume three colleges with payoffs $u_1 = 1$, $u_2 = u < 1$, $\alpha_1 = \alpha_2 = \alpha$, $u_3 = v < u$, $\alpha_3 = \alpha' > \alpha$, and $\alpha u > \alpha' v$. Suppose that the student is either accepted in both 1 and 2 (chance α) or rejected in both. Exhaustive checking reveals that $\Sigma_2(2) = \{1, 3\}$, whereas $Z_2 = \{1, 2\}$.

B. Portfolio choices are less aggressive than sequential choices

Consider the case where a student can apply to the colleges sequentially, observing whether one accepts her before she applies to the next. For a fair comparison, let us restrict to constant marginal costs $c(|S|) = \bar{c}|S|$, $\bar{c} > 0$. The optimal policy in Weitzman (1979) is derived as follows. To each college i, associate an intrinsic *index* or reservation value I_i; this leaves the student indifferent between a final payoff I_i, and first applying to college i and then earning payoff I_i if rejected. Then $I_i = z_i - \bar{c} + (1 - \alpha_i)I_i$ and thus $I_i = (z_i - \bar{c})/\alpha_i = u_i - \bar{c}/\alpha_i$. The optimal policy orders the colleges by their indices I_i; the student proceeds down the list, stopping when a college accepts him (because $u_i > I_i$).

The solution of our static problem substantially differs from the sequential approach. For instance, we have shown that one must apply to the college with the largest expected payoff z_i. Easily, this need not coincide with the college that has the highest Gittins index I_i.

In general, the sequential decision maker employs a more aggressive strategy than does our optimal portfolio strategy. Let W be the list of colleges for which it is sequentially optimal to search, given continued failure, and let $W_{|\Sigma^*|}$ be the set with the $|\Sigma^*|$ highest indices I_i.

THEOREM 3: *The best portfolio Σ^* is not larger than W, and is less aggressive than $W_{|\Sigma^*|}$.*

PROOF: For the size comparison, consider that the sequential rule continues as long as $I_i \geq 0$ or $z_i \geq \bar{c}$. The static decision maker, by contrast, stops when the marginal value of the last college i—which is at most $z_i - \bar{c}$, due to the externalities—turns negative.

We now show that $W_{|\Sigma^*|} \succeq \Sigma^*$. It suffices to show that if $i < j$ and S is *any* portfolio for which MIA picks i over j, then the Gittins indices are likewise ranked $I_i > I_j$. This is obvious if $\alpha_i > \alpha_j$, because $I_i = u_i - \bar{c}/\alpha_i > u_j - \bar{c}/\alpha_j = I_j$. Otherwise, using the marginal benefit expression $MB_k(S) = \rho(U_k)[z_k - \alpha_k f(L_k)] > 0$ from (6), we find that

$$I_i - I_j = \frac{z_i - \bar{c}}{\alpha_i} - \frac{z_j - \bar{c}}{\alpha_j}$$

$$= \frac{1}{\alpha_i}\left(\frac{MB_i(S)}{\rho(U_i)} - \bar{c}\right) - \frac{1}{\alpha_j}\left(\frac{MB_j(S)}{\rho(U_j)} - \bar{c}\right) + [f(L_i) - f(L_j)].$$

If $MB_i(S) \geq MB_j(S)$, then $I_i > I_j$ because $\rho(U_i) < \rho(U_j)$, $\alpha_i < \alpha_j$, and $f(L_i) > f(L_j)$.   *Q.E.D.*

To see that the order can be strict, assume three colleges, again with $\alpha_1 = 0.1$, $\alpha_2 = 0.9$, $\alpha_3 = 1$, $u_1 = 1$, $u_2 = 0.5$, and $u_3 = 0.48$, but now $\bar{c} = 0.05$. One can show that $W_{|\Sigma^*|} = \{1, 2\}$, which is strictly more aggressive than $\Sigma^* = \{1, 3\}$.

5.2. *Portfolio Choice Sets Are Upwardly Diverse*

We turn to another key characteristic of the statically optimal set. How similar should the chosen options be? Is the optimal set an "interval," say $[i, j]$?

Assume first $z_1 > z_2 > \cdots > z_n$. It follows from Theorem 2 that the students should just apply to an interval of top schools. Indeed, $\Sigma^* \succeq Z_{|\Sigma^*|} = [1, |\Sigma^*|]$ implies $\Sigma^* = [1, |\Sigma^*|]$.

Apart from this case, a force to gamble upward emerges and the optimal portfolio is not, in general, an interval. To see this, consider a stylized world with constant marginal cost $\bar{c} > 0$, one college i, and $N - 1$ copies of college $j > i$ with $z_j > z_i$. The algorithm starts with j. We claim that for N large enough and $\bar{c}$ small enough, the algorithm chooses college i before exhausting college j copies. Indeed, suppose the algorithm has chosen $n - 1$ copies of college j, but not yet college i. The marginal benefit of choosing another college j copy is $(1 - \alpha_j)^{n-1}\alpha_j u_j - \bar{c}$. While this vanishes geometrically fast in n, the marginal benefit of choosing college i, namely $\alpha_i u_i - \alpha_i u_j(1 - (1 - \alpha_j)^{n-1}) - \bar{c}$, tends to $\alpha_i(u_i - u_j) - \bar{c}$. For small $\bar{c}$, this is positive. Thus, for large n, it is optimal to choose i over another copy of j. By continuity, this result obviously holds even when all copies of j are not literally identical and there is a sufficiently dense and diverse collection of colleges. So for low enough application costs, one always has an incentive to gamble upward and apply to a discretely higher college than the rest.

6. COMPARATIVE STATICS

We finally consider some natural comparative statics. Obviously, with greater costs $c(\cdot)$, the size of Σ^* decreases, because the algorithm stops sooner.

More interestingly, how will choices adjust when acceptance chances $(\alpha_1, \ldots, \alpha_N)$ change? The answer is far from obvious, because the submodular character of f precludes any direct application of the monotone comparative statics results (see Topkis (1998)).

THEOREM 4: *Assume $\beta = (\beta_1, \ldots, \beta_N)$ is higher than α and proportionately favors better options more than α. Namely, $\beta_i \geq \alpha_i$ for all i and $\beta_i/\alpha_i > \beta_{i+1}/\alpha_{i+1}$, for all $i < N$:*

(a) *The best n-portfolio Σ_n^β is more aggressive than Σ_n^α, or $\Sigma_n^\beta \succeq \Sigma_n^\alpha$.*

(b) *Let $z_i^\alpha = \alpha_i u_i$ and $z_i^\beta = \beta_i u_i$ for all i, and let $z_1^\alpha > \cdots > z_N^\alpha$ and $z_1^\beta > \cdots > z_N^\beta$. The optimal sets are thus $[1, n^\alpha]$ and $[1, n^\beta]$. If also $(1 - \alpha_1)\alpha_2 > (1 - \beta_1)\beta_2$, then $n^\beta \leq n^\alpha$.*

PROOF OF (a): The proof is another double induction on n and N. Let $\Sigma_n^\alpha(N)$ be the optimal n-choice set from N for acceptance chances α. The result holds for $n = 1$ and all N by the MIA. Otherwise, if $j = \arg\max_i \beta_i u_i > \arg\max_i \alpha_i u_i = k$, then $\beta_j u_j \geq \beta_k u_k$ and $\alpha_k u_k \geq \alpha_j u_j$ imply $\beta_j/\beta_k \geq u_k/u_j \geq \alpha_j/\alpha_k$, contrary to our premise.

Assume the result holds for all $\hat{n} \leq n$ and $\hat{N} \leq N$, with one inequality strict. If some $j \notin \Sigma_n^\alpha(N) \cup \Sigma_n^\beta(N)$, then the result holds by induction on the domain $N - j$. Assume there are no omitted options j. Thus, $1 \in \Sigma_n^\alpha(N) \cup \Sigma_n^\beta(N)$. If $1 \in \Sigma_n^\alpha(N) \cap \Sigma_n^\beta(N)$, then

$$\Sigma_n^\beta(N) = 1 + \Sigma_n^\beta(N - 1) \succeq 1 + \Sigma_n^\alpha(N - 1) = \Sigma_n^\alpha(N)$$

by Lemma 1. If $1 \notin \Sigma_n^\alpha(N)$, pick the least $k \notin \Sigma_n^\beta(N)$. Putting $M = [2, k - 1]$, we have

$$\Sigma_n^\beta(N) = 1 + M + L^\beta \succeq M + k + L^\alpha = \Sigma_n^\alpha(N),$$

where $L^\alpha = \Sigma_n^\alpha(N) \cap [k, N]$ and $L^\beta \equiv \Sigma_n^\beta(N) \cap [k, N]$ by Lemma 1. Because $|L^\alpha| = |L^\beta|$, we have $L^\beta \succeq L^\alpha$ by the induction assumption. If $1 \notin \Sigma_n^\beta(N)$, then we can likewise decompose

$$\Sigma_n^\alpha(N) = 1 + M + L^\alpha \quad \text{and} \quad \Sigma_n^\beta(N) = M + k + L^\beta,$$

where $L^\beta \succeq L^\alpha$. The Appendix proves $f^\beta(1 + M + L^\beta) > f^\beta(\Sigma_n^\beta(N))$, contradicting $\Sigma^\beta(N)$ optimal. This case cannot therefore arise. *Q.E.D.*

7. CONCLUSION

Static optimization is rapidly becoming yesterday's struggle in economics. In this paper, we have identified a common and yet unsolved class of downward-recursive static portfolio choice problems—typically, where one earns only the best prize from a portfolio. Such portfolio choices are intriguing, insofar as the *value of a portfolio is subtly less than the sum of its parts.* Such problems are also practically important, being faced by millions of college applicants, thousands of employers competing to hire in student-driven job markets, as well as firms that choose among uncertain technologies to explore.

We have shown that a greedy algorithm finds the optimal portfolio and we have identified the key properties that account for its success. This defines a useful class of submodular functions that can be efficiently maximized. We have also provided some interesting properties that the optimal set possesses.

Chade and Smith (2005) proved that the MIA also works with non-DR functions on a richer set of prizes that satisfy a "second-order stochastic dominance" condition. Examples of the failure of the MIA are given, such as with different option costs or general binary prizes. It is an exciting open problem to find an algorithm that works efficiently in these cases: future research beckons.

Dept. of Economics, Arizona State University, Tempe, AZ 85287-3806, U.S.A.; hector.chade@asu.edu

and

Dept. of Economics, University of Michigan, Ann Arbor, MI 48109-1220, U.S.A.; lones@umich.edu, www.umich.edu/~lones.

APPENDIX: PROOF OF THEOREM 4 FINISHED

Part (a). We need $f^\beta(1 + M + L^\beta) > f^\beta(\Sigma_n^\beta(N))$. If $\alpha_1 > \alpha_k$, then $z_1 > z_k$ and the claim follows from Lemma 2. Assume hereafter $\alpha_1 < \alpha_k$. Then $f^\alpha(\Sigma_n^\alpha(N)) \geq f^\alpha(M + k + L^\alpha)$, because $\Sigma_n^\alpha(N)$ is optimal for f^α. Hence,

$$\alpha_1 u_1 + (1 - \alpha_1)[f^\alpha(M) + \rho^\alpha(M)f^\alpha(L^\alpha)]$$
$$\geq f^\alpha(M) + \rho^\alpha(M)[\alpha_k u_k + (1 - \alpha_k)f^\alpha(L^\alpha)].$$

This holds if and only if

$$(7) \qquad \frac{\alpha_1}{\alpha_k}\left(\frac{u_1 - f^\alpha(M)}{\rho^\alpha(M)}\right) + \left(1 - \frac{\alpha_1}{\alpha_k}\right)f^\alpha(L^\alpha) \geq u_k.$$

We now argue that replacing α by β yields a strict inequality in (7), which is likewise equivalent to $f^\beta(1 + M + L^\beta) > f^\beta(\Sigma_n^\beta(N))$. We now justify this assertion:

• Because 1 dominates every option in $M + L^\alpha$ and $M + L^\beta$, we have $u_1 > f^\alpha(M + L^\alpha)$ and $u_1 > f^\beta(M + L^\beta)$. Using (1), these are equivalent to

$$\frac{u_1 - f^\alpha(M)}{\rho^\alpha(M)} > f^\alpha(L^\alpha) \quad \text{and} \quad \frac{u_1 - f^\beta(M)}{\rho^\beta(M)} > f^\beta(L^\beta).$$

• Since $\beta_1/\beta_k > \alpha_1/\alpha_k$, the weight on the first term of (7) strictly increases.
• We have $f^\beta(L^\beta) \geq f^\beta(L^\alpha) > f^\alpha(L^\alpha)$, respectively, by Lemma 1, and because $\beta_i \geq \alpha_i$ for all i and $\beta_i > \alpha_i$ for some i (because the ratio ordering $\beta_i/\alpha_i > \beta_{i+1}/\alpha_{i+1}$ is strict).
• Finally, the first term in (7) increases as well, because

$$(8) \qquad \frac{\partial}{\partial \alpha_\ell}\left(\frac{u_1 - f^\alpha(M)}{\rho^\alpha(M)}\right) > 0 \quad \forall \ell \in M$$

$$\Rightarrow \quad \frac{u_1 - f^\beta(M)}{\rho^\beta(M)} > \frac{u_1 - f^\alpha(M)}{\rho^\alpha(M)}.$$

To see (8), write $f^\alpha(M) = f^\alpha(U) + \rho^\alpha(U)[\alpha_\ell u_\ell + (1 - \alpha_\ell)f^\alpha(L)]$ using (1), where $L = (\ell, N] \cap M$ and $U = [1, \ell) \cap M$. Thus,

$$
\begin{aligned}
u_1 - f^\alpha(M) &= u_1 - \big(f^\alpha(U) + \rho^\alpha(U)[\alpha_\ell u_\ell + (1 - \alpha_\ell)f^\alpha(L)]\big) \\
&= [u_1 - f^\alpha(U) - \rho^\alpha(U)f^\alpha(L)] - \rho^\alpha(U)[u_\ell - f^\alpha(L)]\alpha_\ell \\
&\equiv A - B\alpha_\ell,
\end{aligned}
$$

thereby implicitly defining A and B. We have

$$
\frac{\partial}{\partial \alpha_\ell} \frac{A - B\alpha_\ell}{\rho^\alpha(M - \ell)(1 - \alpha_\ell)} = \frac{A - B}{\rho^\alpha(M - \ell)(1 - \alpha_\ell)^2}
$$

because $\rho^\alpha(M) = \rho^\alpha(M - \ell)(1 - \alpha_\ell)$. However, this is positive given

$$
\begin{aligned}
A - B &= [u_1 - f^\alpha(U) - \rho^\alpha(U)f^\alpha(L)] - \rho^\alpha(U)[u_\ell - f^\alpha(L)] \\
&= u_1 - f^\alpha(U) - \rho^\alpha(U)u_\ell > 0
\end{aligned}
$$

because option 1 strictly dominates options in $[1, \ell] \cap M$.

Part (b). Parameterize $\theta = \alpha, \beta$, where $\alpha = \theta_L$ and $\beta = \theta_H$. When $z_1^\theta > z_2^\theta > \cdots > z_N^\theta$, the restriction to $C = \{S \subseteq N | S = [1, n], n \le N\}$ is without loss of generality. So consider $\max_{S \subseteq C} v(S, \theta)$.

Observe that C is a chain (i.e., a totally ordered set), and hence $v(S, \theta)$ is quasi-supermodular in S. Thus, to show that the maximizer is increasing in θ, we require that the single crossing property holds (Milgrom and Shannon (1994)), namely

$$
v(S_H, \theta_L) - v(S_L, \theta_L) \begin{cases} \ge 0 \\ > 0 \end{cases} \Rightarrow \quad v(S_H, \theta_H) - v(S_L, \theta_H) \begin{cases} \ge 0 \\ > 0 \end{cases},
$$

where $S_H = [1, n_H]$, $S_L = [1, n_L]$, and $n_L > n_H$. Rewrite these as

$$
\begin{aligned}
\rho^{\theta_L}([1, n_H])f^{\theta_L}((n_H, n_L]) &\begin{cases} \le \\ < \end{cases} c(n_L) - c(n_H) \\
\Rightarrow \quad \rho^{\theta_H}([1, n_H])f^{\theta_H}((n_H, n_L]) &\begin{cases} \le \\ < \end{cases} c(n_L) - c(n_H)
\end{aligned}
$$

for which a sufficient condition is

$$
\rho^{\theta_H}([1, n_H])f^{\theta_H}((n_H, n_L]) < \rho^{\theta_L}([1, n_H])f^{\theta_L}((n_H, n_L]).
$$

Algebra reveals that this holds if $\beta_2(1 - \beta_1) \le \alpha_2(1 - \alpha_1)$ and $\beta_i/\alpha_i > \beta_{i+1}/\alpha_{i+1}$.   *Q.E.D.*

REFERENCES

ALBRECHT, J., P. GAUTIER, AND S. VROMAN (2002): "Equilibrium Directed Search with Multiple Applications," Mimeo, Georgetown. [1294]

BURDETT, K., AND K. JUDD (1983): "Equilibrium Price Dispersion," *Econometrica*, 51, 955–969. [1294]

BURDETT, K., S. SHI, AND R. WRIGHT (2001): "Pricing and Matching with Frictions," *Journal of Political Economy*, 109, 1060–1085. [1294]

CHADE, H., AND L. SMITH (2005): "Simultaneous Search," Mimeo, Michigan and Arizona State. [1294,1301,1304]

GUL, F., AND E. STACCHETTI (1999): "Walrasian Equilibrium with Gross Substitutes," *Journal of Economic Theory*, 87, 95–124. [1300]

KELSO, A. S. J., AND V. P. CRAWFORD (1982): "Job Matching, Coalition Formation, and Gross Substitutes," *Econometrica*, 50, 1483–1504. [1300]

LOVÁSZ, L. (1982): "Submodular Functions and Convexity," in *Mathematical Programming: The State of the Art*, ed. by A. Bachem, M. Grotschel, and B. Korte. Berlin: Springer-Verlag, 235–257. [1300]

MILGROM, P., AND C. SHANNON (1994): "Monotone Comparative Statistics," *Econometrica*, 62, 157–180. [1306]

MUROTA, K., AND A. SHIOURA (2003): "Quasi M-Convex and L-Convex Functions—Quasi-convexity in Discrete Optimization," *Discrete Applied Mathematics*, 131, 467–494. [1301]

STIGLER, G. (1961): "The Economics of Information," *Journal of Political Economy*, 69, 213–225. [1293]

TOPKIS, D. (1998): *Supermodularity and Complementarity*. Princeton, NJ: Princeton University Press. [1303]

WEITZMAN, M. (1979): "Optimal Search for the Best Alternative," *Econometrica*, 47, 641–654. [1293,1302]