Simultaneous search, reservation fees, and sequential outcomes

Ying-Ju Chen
School of Business and Management (ISOM),
The Hong Kong University of Science and Technology,
Clear Water Bay, Kowloon, Hong Kong;
tel: +852-2358-7758; fax: +852-2358-2421;
email: imchen@ust.hk (Corresponding author)

April 2025

Abstract

This paper studies a simultaneous-search problem in which a player observes the outcomes sequentially, and must pay reservation fees to maintain eligibility for recalling the earlier offers. We use postgraduate program applications to illustrate the key ingredients of this family of problems. We develop a parsimonious model with two categories of schools: reach schools, which the player feels very happy upon joining, but the chance of getting into one is low; and safety schools, which are a safer choice but not as exciting. The player first decides on the application portfolio, and then the outcomes from the schools applied to arrive randomly over time. We start with the extreme case wherein the safety schools always admit the player, and show that it suffices to focus on the last safety school. This allows us to conveniently represent the player's value function by a product form of the probability of entering the last safety period and the expected payoff from then on.

We show that the player's payoff after applications is increasing and discrete concave in both the numbers of reach and safety schools, and the optimal number of reach schools increases in the reservation fee. The proof technique utilizes stochastic coupling, stochastic dominance, and directional monotone comparative statics arguments. We also develop a recursive dynamic programming algorithm when admissions to safety schools are no longer guaranteed. We demonstrate instances in which the player applies to more safety schools when either the reservation fee gets higher or the admission probability drops lower, and articulate how these arise from the portfolio optimization consideration.

Keywords: simultaneous search, dynamic programming, endogenous imperfect recall, directional monotone comparative statics

Received: September 2024; accepted: April 2025 by Michael Pinedo after one revision

1 Introduction

Life-changing decisions can occasionally take a long time to make. People may spend months or years preparing for events such as getting licenses, taking exams, planning to get strong letters of recommendation, practicing "LeetCode", and polishing CVs before they battle for industry jobs, academic positions, or higher education programs. After getting themselves ready, people choose the places they apply to, and it is common to see synchronized "seasons" for these submissions. Academic job markets may require candidates to submit the packages before annual conferences, recent graduates must attend campus career fairs for initial screening, and graduate program applicants must collect all the recommendation letters, certified transcripts and standard test scores prior to the admission deadlines. These synchronizations are made for ease of consolidating all the information for the committees to evaluate. Submissions are, however, not the end of the story. Waiting, praying, and getting heart-breaking rejections are commonplace. Even receiving an offer of admission may not necessarily be good news, since the candidates then must make tough decisions regarding accepting it or declining and when and how to conclude the entire journey.

In these examples, decision makers hardly have any bargaining power to push for personalized applications and, therefore, must adapt to the common recruitment schedules. After submitting the applications, the revealed outcomes are random, and so are the moments when these outcomes are revealed. Recruitment committees may shortlist candidates or pursue them with a priority in mind, while applicants need to wait patiently. At the heart of these scenarios is a **two-stage** decision making process. The decision maker (*player* hereafter) needs to first choose from among the set of available options her intended choices (wishlist). Afterward, these intended choices will reveal whether there is a match with the player, typically dynamically; the player then needs to decide whether to and which one to take among the feasible options.

Stage 1 is a *portfolio optimization* problem: the player balances between her dream places and options that are more likely to be a match. Stage 2 is similar to, but more involved than, an optimal stopping time problem. The outcomes are uncertain, and some intended options reveal theirs before others. Thus, the player needs to determine whether to take up the currently available ones, or to keep looking in pursuit of her dream. We notice that a common practice from postgraduate programs is to give a reasonable but short deadline, and charge the *reservation fee* if the candidates indicate their willingness to join, without the binding obligation to eventually register. The payment is non-refundable and must be made in full before the due date. For example, some Hong Kong taught postgraduate (Master of Science or Master of Engineering) programs vary from $70K - $170K, which are quite substantial compared with program tuition fees $200K-$310K.[1] The option to pay reservation fees to maintain eligibility for recalling the earlier offers constitutes the departure from the classical optimal stopping time problem. Thus, the set of feasible options in stage 2 is *endogenously* determined, and may evolve over time depending on the past history of outcome realizations.

As search theory has been a central topic in the economics and marketing literature (Armstrong (2017), Chade et al. (2017), and Rogerson et al. (2005)), to our knowledge no prior work studies exactly the same problem that features simultaneous search and sequential outcomes (the detailed discussions of our paper positioning and its connection to the literature are relegated to Section 2). We use postgraduate program applications to illustrate the key ingredients of this family of problems: simultaneous search, sequential outcomes, and reservation fees are required to maintain the eligibility for recalling the earlier offers. We develop a parsimonious model with two categories of schools: **reach schools** and **safety schools**, each of which has an infinite number of schools. A single applicant (player) obtains a higher payoff upon entering a reach school, but the chance of getting into one is low. In contrast, a safety school is a safer choice but not as exciting, but it is better than leaving empty-handed. All schools within a category are fully interchangeable. In stage one, the player decides how many reach schools and how many safety schools to apply to and pays the application fees.

Stage two is decomposed into periods, and in each period, exactly one school reveals whether her application is successful. Importantly, the sequence in which schools reveal her application outcome is **uniformly random** among the pool she applied to. If her application to a school is successful, she is then asked to decide on the spot whether to (1) accept the offer immediately; or (2) pay the (lump-sum) reservation fee to keep the admission. At the end of stage two, she must decide which school to attend if her consideration set is nonempty. The uniformly random realization setting mimics the rolling-basis admission decisions, which most postgraduate school programs adopt (except US admissions with scholarships, i.e., the April 15 resolution).[2] We formulate this application optimization problem and solve for the optimal selection and stopping rules for this search model.

We start with the extreme case in which the safety schools always admit the player. We show that it suffices to focus on the last safety school, as she will forego all other safety

[1] As these figures in 2022 are jaw-dropping in the eyes of applicants, it even made the headline of popular social media. See `https://mp.weixin.qq.com/s/hiexuamwRY52ZRs60w5baw`. Moreover, in just 3 years (2021 - 2024), some programs have increased their reservation fees by 36.8% (`https://www.nanxingjiaoyu.com/news/5569.html`)! (accessed 2025/03/01)

[2] The timing for sending out rolling-basis admissions may depend on committees' meeting schedules, and how they split the batches of qualified candidates. For the April 15 resolution, see `https://cgsnet.org/resources/for-current-prospective-graduate-students/april-15-resolution/`. (accessed 2025/03/01)

schools given the sure admission later. This observation allows us to conveniently represent the player's value function by a product form of probability of entering the last safety period and the expected payoff from then on, where the latter is independent of the number of safety schools applied to. Exploiting this separation and payoff monotonicity among different last safety school periods, we show that excluding the application fees for safety schools, having more safety schools in the portfolio is always better. Moreover, the marginal benefit of adding more reach or safety schools is decreasing. In all scenarios, the only probability of paying the reservation fee is for the last safety school she encounters. We identify the necessary and sufficient condition for the player to apply for safety schools with no reservation fee (i.e., perfect recall). In addition, we show that if no safety school is applied under perfect recall, the player will focus exclusively on reach school applications when a positive reservation fee is imposed. We also show the optimal number of reach schools is increasing in the reservation fee.

We extend our analysis to accommodate the general scenario such that the admissions of safety schools are not guaranteed. In this case, it is no longer true that the player never accepts a safety school's admission prior to the last safety school period. However, once the player pays the reservation fee for admission to one safety school, she will ignore all subsequent safety schools, irrespective of their application outcomes. This step is critical in devising a two-step dynamic programming algorithm. We demonstrate instances in which the player applies to more safety schools when either the reservation fee is higher or the admission probability becomes lower and articulate how these arise from the portfolio optimization consideration. This has strong managerial implications for service providers in devising their reservation fees and deposits, especially for those institutions that are not universally favored by prospective applicants.

The rest of this paper is organized as follows. Section 2 reviews the relevant literature. Section 3 lays out the key modeling ingredients. In Section 4, we focus on the case wherein the safety schools' admissions are guaranteed, and Section 5 extends the analysis to accommodate uncertain safety schools. Section 6 examines a few variants of model setups, including correlated outcomes among schools, deposits rather than reservation fees, nonuniform randomization of outcome revelations, and heterogeneous reservation fees. We offer some concluding remarks in Section 7, and relegate the technical proofs to the appendix.

2 Literature review

Our study belongs to the long standing literature on search theory. As described in Chade et al. (2017), the literature is divided into two broad categories: simultaneous search and sequential search. In a simultaneous search model, the player decides how many options or which options to search at one time, and chooses one option after seeing all their realizations, as if all the outcomes are revealed simultaneously. This imitates the situation wherein it takes long time for the search preparation (such as school applications, job hunting, or marriage market engagement), and thus the time frame does not allow the player to revise her search strategy conditional on observing some search realizations. In addition, the player is patient enough to wait for all the outcome realizations to finalize her decision, or equivalently, the searched options are not empowered to push the player to decide earlier (such as college applications or the April 15 resolution in the US). See the seminal paper by Stigler (1961) and some recent work by Ali and Shorrer (2025), Chade and Smith (2006), and Smith (2006).

In a sequential search model, the search preparation and outcome realization are both in a short time frame; thus, the player can decide which option to search now, observe the outcome, and then decide whether to keep looking for the next. That is, the player has full control over her selection and stopping. This framework was investigated in detail in the seminal work by Weitzman (1979) and followed up by, e.g., Carrasco and Smith (2017) and Doval (2018). Daughety et al. (1992) interpret options as products offered by different firms; these firms endogenously set limited time windows for recall and the player cannot influence whether an option is recallable. Our problem lies in between in that the player must comply with the schools' pre-specified application window (typically synchronized before the Winter break). Nonetheless, the application outcomes are rolled out over time by the various admission committees during Spring term, and the player has no bargaining power to put the offers on hold without costs.

Because the player can choose the set of schools to apply for upfront, this problem is also related to directed/ targeted search literature (Choi et al. (2018), Eeckhout and Kircher (2010), and Petrikaitė (2018)). Nonetheless, it diverts from this literature in that the outcomes are revealed randomly in the ex post stage. Since the sequence of outcomes goes beyond the player's full control, the problem also has the flavor of random search (Burdett and Vishwanath (1988) and the survey by Rogerson et al. (2005)). Imperfect recall via the channel of reservation fees is less explored, if not completely unexplored, in these research streams. Uniform randomization in conjunction with costly recall creates interactive effects between reach and safety schools. Therefore, even though there is no cost synergy or exogenously imposed application quota (see, e.g., Ali and Shorrer (2025) and Chade and Smith (2006)), the player will take into consideration the numbers of both reach and safety schools altogether despite the separate application costs.

The search theoretic framework has recently been brought to the operations community. Liu et al. (2019) consider a seller with two substitute products who can design the information provision policy to consumers. The accessibility of product information leads to different consumer search behaviors and seller profitability. Wu et al. (2022) investigate the bundling decision of two products when they may have differential information accessibility. The seller uses the operational lever – bundling to manage consumer search behaviors. Because products can be bundled for sale, consumers may purchase products without prior search, thereby leading to a natural application of "nonobligatory search." Hu and Xiao (2024) consider the sequential search problem with ex ante homogeneous boxes and introduce a general recall function among the searched boxes. The seller does not observe the customer's valuation towards its own product and the search outcomes, and decides the time-varying offer. While their setting allows for very general imperfect recall, in our paper recall is endogenously determined by the player's optimal choice of reservation; furthermore, once reserved the valuation from recalling the earlier offer is fixed. Our paper adds to this exciting agenda: embedded in our simultaneous-search setting is an elegant optimal stopping time problem that has "long been explored in operations research" (Chade et al. (2017)).

On the methodological side, our investigation of discrete concavity is relevant to the literature on discrete convexity in the inventory literature. The most celebrated properties are $L^{\natural}-$convexity and $M^{\natural}-$convexity, which are widely adopted for structural properties in inventory management, appointment scheduling, dock allocation, portfolio contract, multi-sourcing, and assemble-to-order systems (see Chen and Li (2021a) and Chen and Li (2021b) for comprehensive surveys). In many known applications with dynamic decision making, the convexity is established to ensure that *per-period* decisions are well-behaved (such as production quantity and exercise amount). However, the number of safety schools is a one-off decision that affects the subsequent dynamic problem of accepting, declining, and reserving.

Our reservation fee is reminiscent of portfolio contract (or capacity reservation) in the supply chain management problem, wherein a retailer or a newsvendor must pay a fee upfront to reserve the right to use some capacity, and after the demand realization, it can decide how much portion to exercise its right (Anderson et al. (2017)). The distinguishing feature is that our "capacity" decisions (the numbers of reach and safety schools) affect the subsequent distributions of realized sequences, and sequencing has substantial impact on the dynamic decision making. As Chen and Li (2021b) point out, the separability ("laminar concavity") structure is key for capacity reservation problems, which does not hold in our setting.

3 Model

We consider the moments when one player applies for schools for postgraduate studies. Suppose that there are two categories of schools: **reach schools** and **safety schools**, each of which has an infinite number of schools. Her payoff upon entering a reach school is u_r, and that upon entering a safety school is u_s, where $u_r > u_s > 0$. However, her application to a reach school is successful only with probability p_r, and her application to a safety school is successful only with probability p_s, where $0 < p_r < p_s \leq 1$. Thus, a reach school is a dream place that she feels very happy about, but the chance of getting into one is low. In contrast, a safety school is a safer choice but not as exciting, and it is better than leaving empty-handed (with payoff 0). These terms are borrowed from the long-standing literature on simultaneous search (see Chade and Smith (2006) for details and Ali and Shorrer (2025) for the college portfolio problem). In our base model, the success probability of each reach school is independent of that of any other reach school, and is independent of that of any safety school. Likewise, the success probability of each safety school is independent of that of any other safety school, and is independent of that of any reach school. In Section 6.1, we analyze an alternative model with correlated outcomes a la Ali and Shorrer (2025).

In stage one, she decides how many reach schools and how many safety schools to apply to. The application fee for a reach school is a_r, and that for a safety school is a_s. In stage two, given the numbers of reach and safety schools N_r and N_s to which she has applied, there are $N_r + N_s$ periods. In each period, exactly one school's outcome is revealed, i.e., whether her application is successful or not, and importantly, the sequence in which schools reveal her application outcome is **uniformly random** among the pool she applied to. Index the school with outcome revealed in period i as school i. If her application to school i is unsuccessful, we move on to period $i+1$.

If her application to school i is successful, she is then asked to decide on the spot whether to (1) accept the offer immediately, in which case the game ends; or (2) pay the (lump-sum) reservation fee F and move on to the next period. If she pays the reservation fee F, she retains the right to join school i at the end, that is, even after passing period i; otherwise she foregoes school i's admission. The homogeneous reservation fee assumption is relaxed in Section 6.4. There is no time discounting, as the player cannot change her postgraduate program commencement date by accepting the admission earlier. In period i, call the set of schools for which she has paid the reservation fee beforehand the consideration set C_i^F. At the end of stage two, she must decide which school to attend if her consideration set is non-empty. The applicant maximizes her expected payoff in each stage, and in stage two maximizes her expected continuation payoff in each period. Figure 1 depicts the sequence of events.

In this problem, the player is allowed to default (i.e., do no show) even if she has paid the reservation fee. Immediate acceptance means she fully commits to a program (like personally calling the coordinator to indicate her acceptance). Also, in some occasions immediate acceptance comes along with the reservation fee, but this will be refunded to her once she joins the school as a tuition deduction. We handle this alternative setup in Section 6.2.[3]

If $F = 0$, the problem degenerates to the classical simultaneous-search one, because the player can always reserve every admission and decide after all outcomes are finalized (as if

[3] More information: https://hkunyouml.hku.hk/content/uploads/2021/06/payment-gk.pdf, https://zhuanlan.zhihu.com/p/362905936, and https://zhuanlan.zhihu.com/p/50570336. (accessed 2025/03/01)

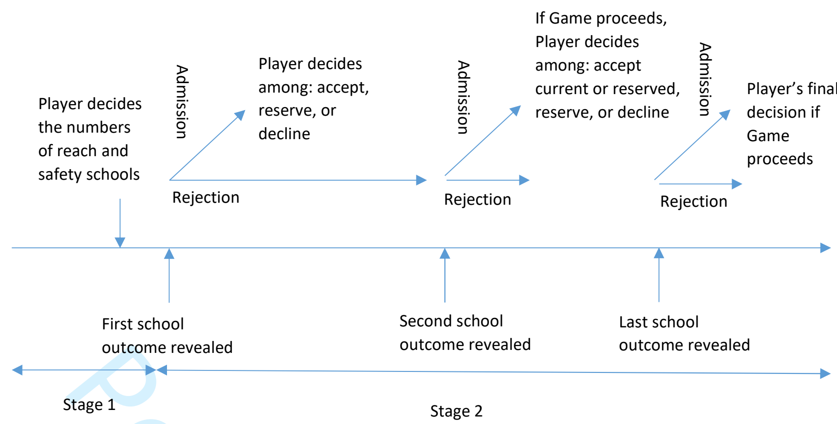

Figure 1: An illustration of timing.

these were simultaneous outcomes). At the other extreme, when $F = \infty$, reservation is infeasible and the player must decide on the spot whether to accept the admission or forego it. The uniformly random realization setting mimics the rolling-basis admission decisions, which most postgraduate school programs adopt (except US admissions with scholarships). If either reach schools reveal their outcomes before any safety school or all safety schools arrive prior to any reach school, the problem becomes less interesting: the player can calculate the probabilities separately and couple these admission probabilities based on the arrival sequence. The general nonuniform randomization scenario is investigated in Section 6.3.

Next, we formulate this application optimization problem and solve for the optimal selection and stopping rules for this search model. To simplify the analysis, we find it helpful to consider the extreme case wherein $p_s = 1$. We will explore this in the next section, and then carry the analysis over to the general case $p_s < 1$.

4 Analysis for $p_s = 1$

This section examines the scenario wherein the safety school always admits the player, i.e., $p_s = 1$. We proceed with the analysis for Stage 2, and then return to the portfolio optimization problem in Stage 1.

4.1 Stage 2; $p_s = 1$

We start with stage 2 – the player's decision after submitting her applications. Our goal is to characterize the player's optimal decision in stage 2, and then derive the player's stage 2 expected payoff. We consider the generic case: $(N_r, N_s) = (R, S)$, where $R \geq 1$, $S \geq 2$, i.e., the player applies to both categories and there are more than one safety schools. First we observe that, as long as there are still some safety schools left, upon seeing a safety school's admission, there is no reason to accept it immediately. This is because the player is guaranteed to be admitted by the last safety school, and waiting for more periods would allow her to see the outcomes of subsequent reach schools. This suggests that when $S \geq 2$, the player will simply **forego the first $S-1$ safety schools.** Let n denote the number of remaining reach schools when the last safety school arrives (the last safety school period). Before hitting this period, in any reach school period, once the player gets an admission, she will accept it and stop immediately. If not admitted by any reach school, the player keeps going and does not pay any reservation fee until the last safety school period.

4.1.1 An extremely imbalanced portfolio

The above discussion suggests that an important building block is the extremely imbalanced portfolio: $(N_r, N_s) = (n-1, 1)$. We find it convenient to count the **number of remaining reach schools** when the unique safety school arrives. Denote this by x, where $x \in \{0, ..., n-1\}$. Accordingly, the number of reach schools that reveal their outcomes before the safety school is $n-1-x$. Within these $n-1-x$ periods, we use n_r to denote the period in which the first reach school issues an admission before the safety school. This occurs with probability $(1-p_r)^{n_r-1}\, p_r$, i.e., the previous reach schools all reject the player. Once a reach school admission comes before the safety school period, the player accepts the offer immediately. Thus, the game ends before the safety school arrives, and the player's payoff is u_r. Summing over all n_r, the probability of accepting a reach school's admission is $1-(1-p_r)^{n-1-x}$, i.e., at least one of these $n-1-x$ reach schools make an offer. With probability $(1-p_r)^{n-1-x}$, the game proceeds to the safety school period when none of these events take place.

Define Π_x as the maximum expected payoff-to-go from the safety school period. The following lemma summarizes the results. In the safety school period, there are three possible strategies upon the safety school admission:

(1) stop in this period;

(2) pay the reservation fee, and keep looking;

(3) decline the offer and do not pay the reservation fee.

Lemma 1. *For $(N_r, N_s) = (n-1, 1)$, let $x \in \{0, ..., n-1\}$ denote the number of remaining reach schools when the unique safety school arrives. The optimal strategy to achieve Π_x is:*

- *If $u_s \geq \left[1-(1-p_r)^x\right] u_r$:*
  - *When $F > \left[1-(1-p_r)^x\right](u_r - u_s)$, choose (1);*
  - *When $F < \left[1-(1-p_r)^x\right](u_r - u_s)$, choose (2);*
- *If $u_s < \left[1-(1-p_r)^x\right] u_r$:*

- *When* $F < (1-p_r)^x\, u_s$, *choose (2);*
- *When* $F > (1-p_r)^x u_s$, *choose (3).*

The player's expected payoff

$$\Pi_x = \max\{u_s, -F + \left[1-(1-p_r)^x\right] u_r + (1-p_r)^x\, u_s, \left[1-(1-p_r)^x\right] u_r\}$$

is weakly increasing and discrete concave in x, *and* $\Pi_x \leq u_r$ *irrespective of the value of* x. *The corresponding expected payoff is*

$$V(n-1,1) = \frac{1}{n}\sum_{x=0}^{n-1}\left(\left(1-(1-p_r)^{n-1-x}\right) u_r + (1-p_r)^{n-1-x}\,\Pi_x\right). \tag{1}$$

We now discuss the results of Lemma 1 in detail. For the safety school period, there are x reach schools left. We need to compare the reservation fee, the probability of receiving an offer of admission by the remaining reach schools, and the incremental utility $u_r - u_s$. The optimal strategy depends on the relative magnitude of these parameters. Furthermore, because the reservation decision depends on x, the consideration set is clearly path-dependent. An early admission from a reach school prior to the safety school will eliminate the consideration set. Even if all prior reach schools reject the player, whether the player reserves the safety school depends on *when* it arrives.

We note that $1-(1-p_r)^x$ is increasing in x. Thus, $u_s > \left[1-(1-p_r)^x\right] u_r$ will hold when there are only a few remaining reach schools. Under this category, $F > \left[1-(1-p_r)^x\right](u_r - u_s)$ is more likely to hold when x is sufficiently small. In this case, the player will accept the offer by the (current) safety school and end the game. As x becomes larger (i.e., more reach schools are left), $F < \left[1-(1-p_r)^x\right](u_r - u_s)$. In this case, the player will pay the reservation fee, and keep looking. When there are more reach schools remaining, the player is more willing to pay F and continue. With the safety school's admission in hand, the player will stop immediately at the first reach school which issues an admission; if no reach school offers any admission, in the end the player will return to the reserved safety school offer.

Now suppose that we have even more reach schools awaiting such that $u_s < \left[1-(1-p_r)^x\right] u_r$. The player will keep looking and need to decide whether to pay F. The option is taken only when all reach schools reject the player. When $F < (1-p_r)^x\, u_s$, the player will pay the reservation fee, and keep looking. Again the player will stop immediately upon receiving the first reach school's admission, and will return to the reserved safety school offer if she eventually runs out of all reach schools. Finally, when there are truly a lot of reach schools such that $F > (1-p_r)^x u_s$, the player will not pay the reservation fee to this safety school. Instead, she bets on the subsequent reach schools. Depending on the parameter values, some of the above regions may degenerate. In Figure 2, we illustrate the corresponding strategies in these regions.

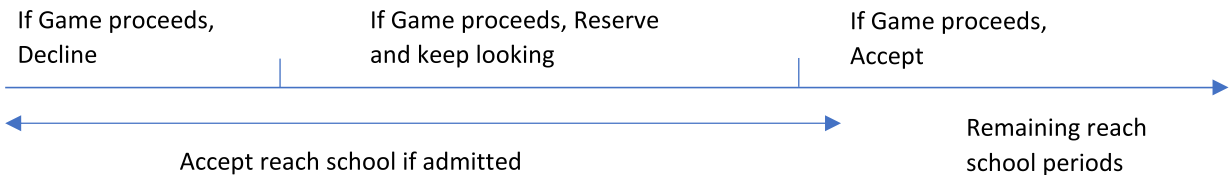

Figure 2: An illustration of optimal strategy in the last safety school period.

Finally, we also make a simple observation that when there are more remaining reach schools after the safety school, the player anticipates a larger expected payoff-to-go. In addition, this benefit gets amplified when fewer and fewer reach schools remain.

4.1.2 Payoff in the generic case

Returning to the general case (R, S), in the last safety school period, the player will follow the strategies depicted in Lemma 1, where we substitute x by n. Put it differently, the environment studied in Lemma 1 can be thought of as the subgame $(N_r, N_s) = (n, 1)$ with the **first period being a safety school period.** Another observation is that following the strategy in Lemma 1, the player's expected payoff will be pinned down by the value of n and is independent of R and S, the numbers of reach and safety schools she starts with.

What is the probability of entering the last safety school period? Among those R reach schools, n of them arrive after the last safety school period. Thus, this event occurs when all those $(R-n)$ reach schools prior to that reject the player, and its probability is $(1-p_r)^{R-n}$. How often would we see that the last safety school period corresponds to n? Some elementary combinatorial calculations are helpful. Before this period, there are $(R-n)$ reach schools and $(S-1)$ safety schools. Thus, the probability, denoted as $\mathbb{P}(n)$, is

$$\mathbb{P}(n) = \begin{cases} 0, & n \geq R+1. \\ \frac{C_{S-1}^{R+S-n-1}}{C_S^{R+S}}, & n = 0, \ldots, R. \end{cases} \tag{2}$$

Note that $\mathbb{P}(n)$ is a well-defined probability mass function for all (R, S); that is, $\sum_{n=0}^{\infty} \mathbb{P}(n) = 1$.

We can also derive the player's expected payoff $V(R, S)$ based on the above:

$$\begin{aligned} V(R,S) &= \sum_{n=0}^{R} \mathbb{P}(n) \left\{ \left[1-(1-p_r)^{R-n}\right] u_r + (1-p_r)^{R-n} \Pi_n \right\} \\ &= \sum_{n=0}^{R} \mathbb{P}(n) \left\{ u_r - (1-p_r)^{R-n} (u_r - \Pi_n) \right\}. \end{aligned} \tag{3}$$

In equation (3), the first term corresponds to the scenario where a reach school's admission comes before the last safety school period, and the second term addresses the alternative scenario.

By Lemma 1, $u_r - \Pi_n \geq 0$, $\forall n$. Moreover, $u_r - \Pi_n$ decreases in n, but $(1-p_r)^{R-n}$ increases in n. To develop the procedure to characterize the two-state dynamic program, we are now ready to explore the property of $V(R,S)$. We show that excluding the application fees, having more reach or safety schools in the portfolio is always better for the player.

Proposition 1. *$V(R,S)$ is increasing in R and S respectively.*

Next, we establish an important property on the marginal benefit for adding more safety schools or more reach schools. This will then imply that the stage 1 problem (of finding the optimal portfolio) is well-behaved. Since the proof techniques are somewhat different, we state them as two separate propositions.

Proposition 2. *For any R, $V(R,S)$ is discrete concave in S.*

In the proof of Proposition 2, we show that

$$V(R,S) - V(R,S-1) = \sum_{n=0}^{R} \{\mathbb{P}(n;R,S) - \mathbb{P}(n;R,S-1)\} \left\{ u_r - (1-p_r)^{R-n}(u_r - \Pi_n) \right\}. \tag{4}$$

This illustrates the rationale of applying for more safety schools. Proposition 2 suggests that if the player could choose when the offer from the safety school arrives, then one would choose it to arrive last. However, the player cannot predict whether reach schools or safety schools arrive earlier. Thus, it is not enough to simply apply just one safety school. Indeed, (4) shows that applying for more safety schools allows the player to increase the chance of such late-arriving safety offers. Mathematically, postponing the last safety school period translates into the first-order stochastic dominance between $\mathbb{P}(n;R,S)$ and $\mathbb{P}(n;R,S-1)$.

Proposition 3. *For any S, $V(R,S)$ is discrete concave in R.*

The proofs of Propositions 1 and 2 utilize heavily (3), where we transform the problem into the probability space discussion. In contrast, to prove discrete concavity with respect to R, we leverage the dynamic programming recursion described in the appendix. As aforementioned in Section 2, the state R enters the probability expectation nonlinearly ($\frac{R}{R+S}$). This creates technical challenges that would not occur in Chen and Li (2021a): when we vary R, both the value-to-go function and the probability distribution are affected in a convoluted manner. Thus, one needs to meticulously decompose their impact via algebraic rearrangement. The proof of Proposition 3 does precisely this to overcome the nonlinearity issue.

4.1.3 Numerical examples

Now we demonstrate some numerical examples below and explore different parameter values. In Figure 2, we have illustrated that the player will pay the reservation fee when the remaining period falls in a non-trivial *range*. Building upon this, Figures 3 – 5 show that this range of paying the reservation fee can exhibit various patterns as we increase the reservation fee F. Common in these figures are that the x-axis is the reservation fee (F), the y-axis is the number of remaining reach schools, and the dark regions correspond to the case in which the player optimally chooses to reserve the last safety school.

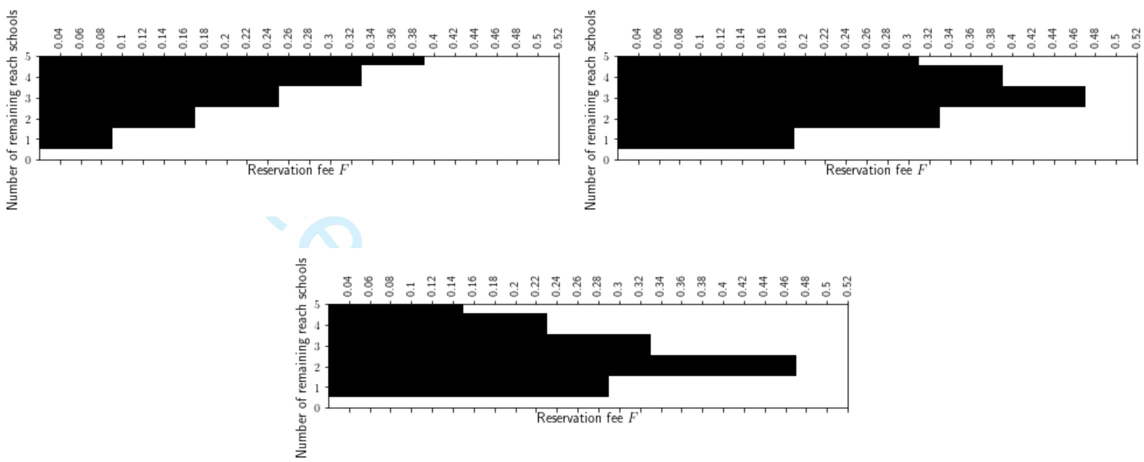

Figure 3: Range of paying the reservation fee. Top Left: $p_r = 0.1$; Top Right: $p_r = 0.2$; Bottom: $p_r = 0.3$. Other parameter values are: $R = 5, S = 2, u_r = 2, u_s = 1$.

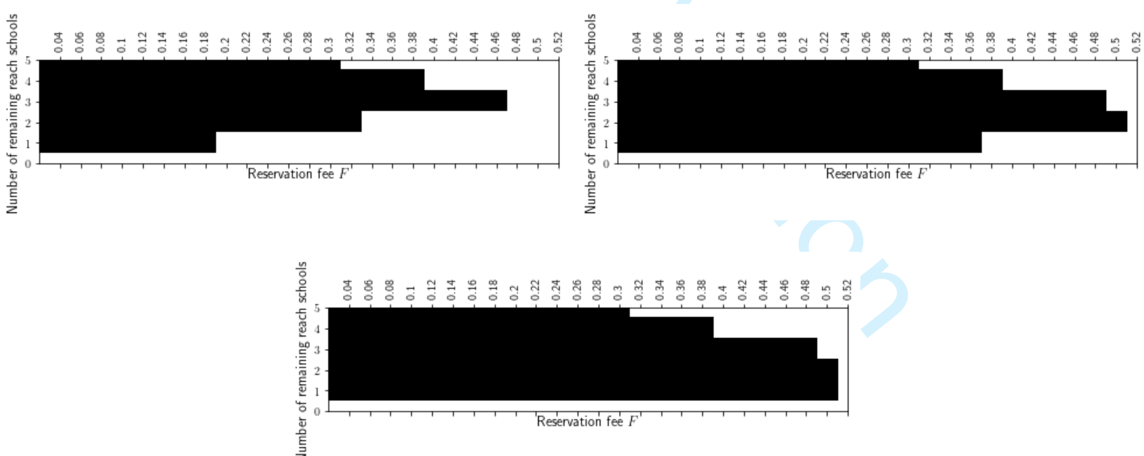

Figure 4: Range of paying the reservation fee. Top Left: $u_r = 2$; Top Right: $u_r = 3$; Bottom: $u_r = 4$. Other parameter values are: $R = 5, S = 2, p_r = 0.2, u_s = 1$.

In the proof of Lemma 1, we note that when n is very large, the player's optimal strategy is to keep looking without paying the reservation fee. At the other extreme, when n is very small, the player will accept the safety school's offer immediately without further search. Thus, in the two ends changing the reservation fee will not influence the player's strategy, nor will this influence the player's expected payoff Π_n. However, in the middle range, when the player prefers to reserve the safety school's admission, a higher reservation fee will hurt the player and reduce Π_n. Figures 3 – 5 suggest that overall, the range shrinks because it becomes more costly to reserve an admission. However, the shrinkage can start from either end or both ends, when we try different values of admission probabilities of reach schools p_r, payoffs of reach schools u_r, and numbers of reach schools R. We note that in Figure 4, the range of paying reservation fees is influenced by u_r. When reach schools become more appealing (u_r gets larger), the player will be more willing to bet on the remaining reach schools. This can be observed from the expression of Π_x in Lemma 1.

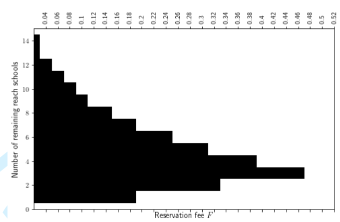

Figure 5: Range of paying the reservation fee. $R = 15, S = 2, p_r = 0.2, u_r = 2, u_s = 1$.

4.2 Stage 1: $p_s = 1$

In this section we return to stage 1 wherein the player decides the portfolio of schools. Below, we start with the simplified case $a_s = 0$, and then examine the general case $a_s > 0$.

4.2.1 Stage 1: $a_s = 0$, $p_s = 1$

When applying for safety school cost $a_s = 0$, we note that the player can always forego all the safety schools except the last one, which suggests that she obtains a weakly higher expected payoff by adding more safety schools (Proposition 1). Hence, for the stage 1, the optimal strategy is to apply for an infinite number of safety schools: $N_s^* = +\infty$, and a finite number of reach schools as it is costly to apply ($a_r > 0$). In this extreme case, the player will not pay any reservation fee to safety schools as there are an infinite number of them. Additionally, the player will not pay any reservation fee to reach schools as she will enter one immediately once admitted.

Given the strategy, her expected payoff is:

$$V(N_r, \infty) = u_s + \left[1 - (1 - p_r)^{N_r}\right](u_r - u_s) - N_r a_r, \tag{5}$$

where $\left[1 - (1 - p_r)^{N_r}\right]$ is the probability that the player ever receives an admission from any reach school, and $N_r a_r$ is the aggregate application fees she pays for the reach schools. Since the player never runs out of safety schools, she is guaranteed to obtain at least the base payoff u_s.

Proposition 4. *Suppose $p_s = 1$ and $a_s = 0$. In stage 1, if $\log_{1-p_r} \frac{a_r}{p_r(u_r - u_s)} < -1$, at optimality $N_r^* = 0$. Otherwise, the optimal N_r^* is given by the integer value that satisfies:*

$$\log_{1-p_r} \frac{a_r}{p_r\left(u_r - u_s\right)} \leq N_r^* \leq \log_{1-p_r} \frac{a_r}{p_r\left(u_r - u_s\right)} + 1.$$

As a remark, the "periods" of stage two in our model mainly indicate the *sequence* of school decisions. Thus, in this limiting case, uniform randomness of school arrivals simply means that each applied school will get equal chances of arriving first, second, etc. Because there is no discount in our model, what really matters is whether there are remaining safety schools for a given decision-making period. Pragmatically, Proposition 4 suggests that the player shall try to apply to many safety schools to minimize the chance of seeing the last safety school while running out of all reach schools. Moreover, even though in equilibrium the player never pays to reserve any school, the existence of reservation fee is key to the rationale of optimal portfolio optimization in stage one. It is precisely the fear of paying unnecessary reservation fees that induces the player to apply infinitely many safety schools, rendering any portfolio with finite schools *strictly suboptimal*.

Insofar we focus on the guaranteed safety school case $p_s = 1$, but it is worthwhile discussing the alternative scenario with non-guaranteed safety schools $p_s < 1$. When $p_s < 1$, although it is not guaranteed that any safety school will admit the player, applying for safety schools is costless. Thus, it remains beneficial for players to include more safety schools in the portfolio, and the optimal strategy is still to apply for an infinite number of safety schools (and finite number of reach schools as it is costly). That is, $N_s^* = +\infty$, $N_r^* < \infty$.

Therefore, the player will enter the first reach school that admits her and pass all the safety schools along the way, without paying any reservation fee to any school. If she runs out of all reach schools, the player will accept the first safety school that offers admission after the last reach school. Given that the player is not admitted by any reach school, the probability of being admitted by a safety school from the last reach school is

$$\sum_{n=0}^{+\infty} \left(1 - p_s\right)^n p_s = p_s \frac{1}{1 - \left(1 - p_s\right)} = 1,$$

where we note that there are still an infinite number of safety schools in any subsequence from then on. Hence, she will be guaranteed to enter a safety school after being rejected by all reach schools, and she will not pay the reservation fee to the safety school. The player's expected payoff under $N_s^* = +\infty$ for $p_s < 1$ has the same expression as equation (5). This gives exactly the same N_r^* as that under $p_s = 1$.

4.2.2 Stage 1: $a_s > 0$, $p_s = 1$

From Stage 2, we know that the key event is the occurrence of the last safety school. Thus, the player has the incentive to postpone this event as late as possible. When $a_s = 0$, the player can freely obtain an infinite number of safety school admission offers and never run out of safety schools. If $a_s > 0$, this becomes costly. Nonetheless, the analytical results established above can be applied to develop handy algorithms to solve the player's dynamic program, and accordingly our structural property also allows us to develop a simple plan for finding the optimal portfolio. While there are many parameters we can perturb to illustrate the usefulness of our framework, we find it more sensible to focus on the new features introduced in this paper, namely the reservation fee F, and the corresponding numbers of reach and safety schools.

When the reservation fee F gets larger, would the player expect to pay more? We offer some answers in Figure 6. In the left panel of Figure 6, we fix the number of safety schools and increase the reservation fee F. Consider for example the case $S = 5$. It is apparent that as F increases, the expected payment for reservation is non-monotone. The three sudden downward jumps occur because the player changes the strategy in the last safety school period. Because the number of remaining reach schools is an integer, a small change in the reservation fee will not trigger the strategy switch. Once the change becomes substantial, the range starts to shrink either from one end or both ends. For a fixed pair of (R, S), the probabilities of entering each scenario Π_n are given by the probabilities $\mathbb{P}(n)$. Therefore, the strategy change will result in a reduction of expected payment for reservation at these switching points. The same phenomenon occurs for the other cases $S = 1$ and $S = 3$.

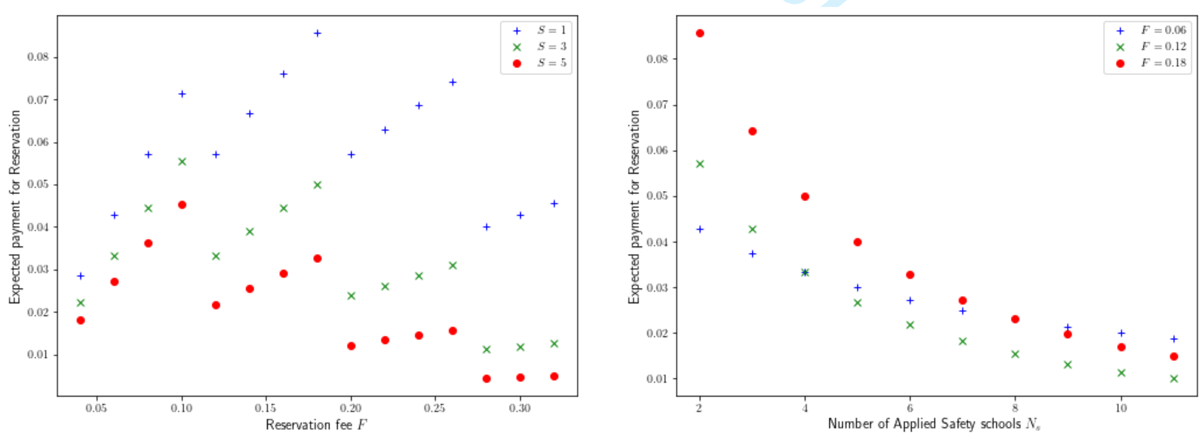

Figure 6: Expected payment for reservation. Other parameter values are: $R = 5, p_r = 0.1, u_r = 2, u_s = 1$.

We also note that, because the strategy induced by maximizing Π_n is independent of the number of applied safety schools, when we vary S, the expected payment for reservation is only influenced by the probabilities $\mathbb{P}(n)$. Consequently, for any given number of remaining schools, a larger S leads to a lower expected payment for reservation because the probabilities are shifted toward scenarios with smaller n. In the right panel of Figure 6, we instead fix the reservation fee F and increase the number of safety schools $N_s := S$. This time we observe that having more safety schools reduces the expected payment for reservation. Facing the reservation fee, the player can apply to more safety schools to push forward the last safety school, and this effectively mitigates the chance of having to pay for reservation.

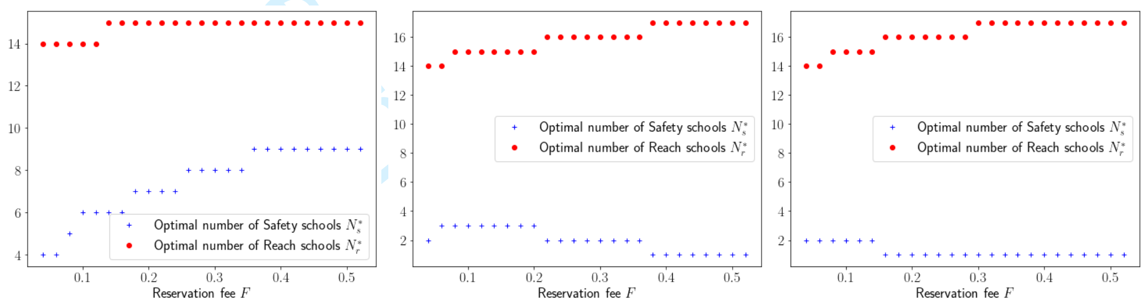

Figure 7: Optimal numbers of reach and safety schools. Left: $a_s = 0.001$; Middle: $a_s = 0.003$; Right: $a_s = 0.005$. Other parameter values are: $a_r = 0.01, p_r = 0.2, u_r = 2, u_s = 1$.

Pushing this idea further, we can examine how the player selects the portfolio in stage 1 when facing different reservation fees. The results are illustrated in Figure 7. Recall that the player pays the reservation fee only for the last safety school she encounters. Thus, one might be tempted to conjecture that when the reservation fee increases, the player shall apply to fewer safety schools. However, this is not always true. As we demonstrate, N_s^* can be increasing, non-monotone, or decreasing in F; whereas in our numerical experiments N_r^* is increasing in F. Let us start with the easy part N_r^*. Recall that the player never pays the reservation fee for any reach school. Thus, intuitively a higher reservation fee enhances the competitive advantage of reach schools over safety schools, and we would expect that the player intends to apply for more reach schools. This is precisely what we observe in Figure 7.

The optimal number of safety schools is more subtle. We note that a higher reservation fee will push the player to either try hard to delay the last safety school, or make it arrive early. When applying to safety schools is not expensive ($a_s > 0$ but small as in Figure 7), delaying the decision stage is more profitable; thus, a higher reservation fee induces the player to apply to more safety schools. This situation is very sensible in reality as in Hong Kong some post-graduate programs charge only \$0.18K HKD for applications but \$70K - \$170K for reserving admissions. On the other hand, uniform randomization over the outcome sequence and imperfect recall create some congestion between reach and safety schools. When the player intends to apply for more reach schools, the room for safety schools gets squeezed. Consequently, a higher reservation fee may drive down the optimal number of safety schools through the increase in reach schools. Taken together, the optimal number of safety schools N_s^* is very sensitive to the application fee a_s, and the pattern can be either increasing, non-monotone, or decreasing in the reservation fee F.

Through the above discussions we also uncover the salient difference between the application fee and the reservation fee. The player may apply for more or fewer safety schools when facing a higher reservation fee in Figure 7. In contrast, a higher application fee unambiguously reduces the number of applied safety schools. To elaborate, the player must pay the application fee before considering any school, but upon admission the player may save the reservation fee and accept it directly. Thus, even though the reservation fee is non-refundable, the option of not paying it with outright acceptance leads to different implications from the application fee. In connection to the search literature, the application fee is the same as the classical search cost, whereas the reservation fee is somewhat similar to the non-obligatory search problem (Doval (2018) and Wu et al. (2022)): this cost can be bypassed at the player's own will.

The above numerical observations lead us to close the loop with an analytical result. As Figure 7 demonstrates that the optimal number of reach schools is increasing in the reservation fee F, but that of safety schools can take arbitrary patterns, the only feasible analytical finding one could hope for is the former. Indeed, the next proposition establishes such monotonicity and provides the analytical ground for our numerical findings.

Proposition 5. *Suppose $p_s = 1$. In stage 1, for any a_r, a_s, the optimal number of reach schools, R^*, monotonically increases in F.*

The proof utilizes a novel mathematical technique: directional monotone comparative statics in the discrete state space. The technical challenges arise in two facets. First, since the optimal number of safety schools can be non-monotone in the reservation fee F, to show Proposition 5, we must establish a structural property: the objective function $V(R, S)$ changes more significantly with respect to F with a higher number of R, for any larger or smaller number of S. This is termed as the directional single-crossing condition. Barthel and Sabarwal (2018) admit this property is *way more demanding* than the usual single-crossing condition (the usual one only requires the verifications when *both* R and S become larger). Second, the literature on such directional monotonicity prior to Barthel and Sabarwal (2018) all assumes some convex sets of the decision variables, but our discrete optimization problem falls outside this space. Therefore, on the methodological front, this project showcases a practical situation in which discrete optimization and directional monotone comparative statics can be analytically handled.

4.2.3 Stage 1: Extreme cases $F = 0$ and $F = \infty$, $p_s = 1$

We can also connect our results to the classic problems where perfect recall is assumed (i.e., $F = 0$). In so doing, we can also comment on the general case with $F > 0$. The school choice papers by Chade and Smith (2006) and Ali and Shorrer (2025) assume that the application cost depends only on the number of applied schools. In other words, this degenerates to a special case with $a_r = a_s > 0$. When $F = 0$, Option (2) in Lemma 1 dominates (1) and (3). Freely reserving the safety school admission allows the player to take up the reserved one when all remaining reach schools issue rejections, and immediately accepting the safety school foregoes the chance of seeing the reach school outcomes. This implies that if in Figure 7 we extend range to cover $F = 0$, N_s^* is either 0 or 1. We now show that

Proposition 6. *Suppose that $a_r = a_s \equiv a > 0$ and $p_s = 1$.*

- *If $F = 0$, the player applies to exactly one safety school if and only if $u_s \geq p_r u_r$.*
- *If the player applies to no safety school under $F = 0$, then the player will not apply to any safety school under any other reservation fee $F > 0$.*

Proposition 6 identifies the necessary and sufficient condition for the player to apply for safety schools with perfect recall ($F = 0$). In addition, it shows that if no safety school is applied under perfect recall, the player will focus exclusively on reach school applications when a positive reservation fee is imposed. The second part applies more generally to other cases (such as non-identical application fees ($a_r \neq a_s$) and non-guaranteed safety school admission $p_s < 1$).

Next, we switch to the other extreme $F = \infty$. When $F = \infty$, the player never pays for reservation, and the player simply compares immediately accepting the safety school with continuing for reach schools. One can then find a simple threshold $\hat{x}$ such that the player accepts the safety school if the remaining number of reach schools is smaller than $\hat{x}$, and rejects it otherwise. In addition to the above $a_r = a_s$ case, Chade and Smith (2006) and Ali and Shorrer (2025) focus on another important case with a fixed number of applications, i.e., $R + S = C$. We hereby examine this scenario and provide some partial results on the comparative statics.

Proposition 7. *Suppose that $F = \infty$ and the player can apply to at most C schools. If $C \leq \log_{1-p_r} \frac{u_r - u_s}{u_r}$, then:*

- *If u_r increases, the player applies to more reach schools and fewer safety schools.*
- *If u_s increases, the player applies to fewer reach schools and more safety schools.*

Proposition 7 confirms our intuition and serves as sanity check that our model is sensible. A higher payoff from a reach school makes it more appealing, and therefore the player applies to more reach schools and fewer safety schools. The converse is true when the player collects a higher payoff from a safety school. In the proof we show that when $C \le \log_{1-p_r} \frac{u_r - u_s}{u_r}$, the expression of $V(R,S)$ takes a simple form. This allows us to apply the first-order approach for comparative statics. If this condition is violated, the optimal decision at period x will depend on the number of remaining reach schools.

5 Analysis for $p_s < 1$

In this section, we extend our analysis to accommodate the alternative scenario $p_s < 1$. When $p_s < 1$, it is not guaranteed that any safety school will admit the player.

5.1 Analytical results: $p_s < 1$

We define the expected payoff-to-go in this case as $V(R,S)$ in general. Similar to the case $p_s = 1$, we start with a simplified scenario of stage 2: $(N_r, N_s) = (n-1, 1)$. We again denote by x the number of remaining reach schools when the unique safety school arrives, where $x \in \{0, ..., n-1\}$. Once a reach school admission comes before the safety school period, the player accepts the offer immediately. With probability $(1-p_r)^{n-1-x}$, the game proceeds to the safety school period, which now generates a *random* admission outcome. In the safety school period, if the outcome is a rejection, the game proceeds to $(N_r, N_s) = (x, 0)$, and the player's expected payoff is $\left[1 - (1-p_r)^x\right] u_r$. If the outcome is an admission, the problem is equivalent to that studied in Section 4.1, and Lemma 1 applies.

Notably, $\tilde{V}(n,S)$ corresponds to the case in which there are n reach schools and S safety schools, *with the first school being a safety one.* In contrast, $V(n, S-1)$ corresponds to the case in which there are n reach schools and $S-1$ safety schools, and the arrival sequence is uniformly random.

Lemma 2. *The corresponding maximum expected payoff for this safety school period, defined by $\tilde{V}(x,1)$, is*

$$\tilde{V}(x,1) = p_s \Pi_x + (1-p_s)\left[1 - (1-p_r)^x\right] u_r,$$

where we recall that Π_x is the expected payoff when $p_s = 1$ in Lemma 1. Moreover, $\tilde{V}(x,1) \le \Pi_x$, and the expected payoff-to-go $V(n-1,1)$ is

$$\begin{aligned} V(n-1,1) &= \sum_{x=0}^{n-1} \frac{C_x^{x+1-1}}{C_1^n} \left\{ \left[1 - (1-p_r)^{n-1-x}\right] u_r + (1-p_r)^{n-1-x} \tilde{V}(x,1) \right\} \\ &= \frac{1}{n} \sum_{x=0}^{n-1} \left\{ u_r - (1-p_r)^{n-1-x} \left[u_r - \tilde{V}(x,1)\right] \right\}. \end{aligned} \tag{6}$$

Lemma 2 demonstrates the tight connection between this section and our base model: if

$p_s = 1$, $\tilde{V}(x,1)$ coincides with Π_x. Now consider the general case: $(N_r, N_s) = (R, S)$, where $R \geq 1$, $S \geq 2$. One might be tempted to follow the procedure for $p_s = 1$ here and express $V(R,S)$ by conditioning on different scenarios of $\tilde{V}(x,1)$ a la equation (3). However, this naive approach does not work. Recall that when $p_s = 1$, as long as there are still some safety schools left, upon seeing a safety school's admission, there is no reason to accept it immediately. This is no longer true because safety school admission is not guaranteed when $p_s < 1$. The expression of $V(R,S)$ in equation (3) actually uses this optimal decision rule implicitly: the probability of entering the scenario n, $\mathbb{P}(n)$, results from the observations that (a) the player will accept a reach school immediately once an admission is issued; and (b) the player will ignore all safety schools' admissions beforehand. When $p_s < 1$, part (a) remains valid because $u_r > u_s$, but part (b) does not apply. Therefore, it is possible that the player may accept or reserve a safety school's admission before hitting "scenario n", i.e., the last safety school arrives with n remaining reach schools.

For this matter, let us consider instead the **first** safety school and suppose that when it arrives there are n remaining reach schools. In other words, the first safety school occurs in the $(n+S)$-th period, counting from the end. Before hitting this period, we know the optimal decision rule: once the player gets an admission in any reach school period, she will accept it and stop immediately. Now imagine that the game proceeds to this first safety school period. With probability $(1-p_s)$, the safety school outcome is a rejection. The game proceeds to the next period and the player faces $(N_r, N_s) = (n, S-1)$. Recall that we define $V(n, S-1)$ as the expected payoff-to-go for this case.

If it is an admission, the player chooses among three options: (1) stop in this period; (2) pay the reservation fee, and keep looking; (3) decline the offer and do not pay the reservation fee. Option (1) gives utility of u_s. For Option (2), **once the player pays the reservation fee, she foregoes all the subsequent** $(S-1)$ **safety schools,** irrespective of their application outcomes. This is because all safety schools yield the same payoff u_s upon joining. Thus, her expected utility is

$$-F + \left[1 - (1-p_r)^n\right] u_r + (1-p_r)^n u_s.$$

For Option (3), she proceeds to the next period with $(N_r, N_s) = (n, S-1)$ and anticipates expected payoff $V(n, S-1)$. Collectively, the player's expected payoff can be written as follows:

$$\tilde{V}(n,S) = (1-p_s)V(n,S-1) + p_s \max\left\{u_s, -F + \left[1-(1-p_r)^n\right] u_r + (1-p_r)^n u_s, V(n,S-1)\right\}. \tag{7}$$

Notice the difference between $\tilde{V}(n,S)$ and $V(n,S-1)$. $\tilde{V}(n,S)$ corresponds to the case in which there are n reach schools and S safety schools, *with the first school being a safety one.* In contrast, $V(n,S-1)$ corresponds to the case in which there are n reach schools and $S-1$ safety schools, and the arrival sequence is uniformly random.

We need to compare the reservation fee, the probability of getting admission by reach schools left, and the incremental utility $u_r - u_s$. The optimal strategy depends on the relative magnitude of these parameters. Finally, the true expected payoff for a player with $(N_r, N_s) = (R, S)$ can be derived from $\tilde{V}(n, S)$, where $n \in (0, \dots, R)$. When $n = R$, all reach schools arrive after the first safety school; this occurs with probability $\frac{S}{R+S} = \frac{C_{S-1}^{R+S-1}}{C_S^{R+S}}$. For general n, the probability is $\frac{C_{S-1}^{n+S-1}}{C_S^{R+S}}$. Recall that the player proceeds to the first safety school period when no prior reach schools admit her, which happens with probability $(1-p_r)^{R-n}$. Consequently, we obtain that

$$\begin{aligned} V(R,S) &= \sum_{n=0}^{R} \frac{C_{S-1}^{n+S-1}}{C_S^{R+S}} \left\{ \left[1-(1-p_r)^{R-n}\right] u_r + (1-p_r)^{R-n}\, \tilde{V}(n,S) \right\} \\ &= \sum_{m=0}^{R} \frac{C_{S-1}^{R+S-m-1}}{C_S^{R+S}} \left\{ \left[1-(1-p_r)^{m}\right] u_r + (1-p_r)^{m}\, \tilde{V}(R-m,S) \right\} \\ &\equiv \sum_{m=0}^{R} \mathbb{P}(m) \left\{ u_r - (1-p_r)^m \left[u_r - \tilde{V}(R-m,S)\right] \right\}, \end{aligned} \tag{8}$$

where $\mathbb{P}(m)$ has already been defined in equation (2).

Here we notice that the expected payoff $V(n, S-1)$ is needed for obtaining $\tilde{V}(n, S)$, where $V(n, S-1)$ depends on all of those $\tilde{V}(i, S-1)$, $\forall i = 0, \dots, n$. Thus, for any n, it will require recursive calculations in order to pin down the values of all expected payoffs $\tilde{V}(n,1), \dots, \tilde{V}(n, S-1)$. We have already known the boundary condition $V(n-1, 1)$ and $\tilde{V}(x, 1), \forall x = 0, \dots, n$, whose expressions are given in equation (6). Following this procedure, we can obtain all expressions of $\tilde{V}(R, S)$ and $V(R, S)$. We formally state the algorithm below.

Algorithm for $V(R,S)$ when $p_s < 1$.

- Step 0. Calculate $V(n, 0) = \left[1-(1-p_r)^n\right] u_r$, $\forall n = 0, \dots, R$.
- Step 1.
  (a) Calculate $\tilde{V}(n, 1), \forall n = 0, \dots, R$, by equation (6).
  (b) Calculate $V(n, 1), \forall n = 0, \dots, R$, by equation (8).
- Step j: j from 2 to S
  (a) Calculate $\tilde{V}(n, j), \forall n = 0, \dots, R$, by equation (7).
  (b) Calculate $V(n, j), \forall n = 0, \dots, R$, by equation (8).

In part (a) of Step j, to obtain $\tilde{V}(n, j)$, we will need $V(n, j-1)$ calculated in part (b) of Step $j-1$; whereas in part (b) of Step j, to obtain $V(n, j)$, we will need $\tilde{V}(m, j), \forall m = 0, \dots, n$ calculated in part (a). Thus, recursively we can fill in all the required values and obtain $V(R, S)$ at the end of Step S. Starting with any state (R, S), we can write down the subsequent development based on the next arrival, which is either a safety school or a reach school.

$$
\begin{aligned}
V(R,S) =& \frac{S}{R+S}\tilde{V}(R,S) + \frac{R}{R+S}\left\{p_r u_r + (1-p_r)V(R-1,S)\right\},\ \forall R \geq 1,\ S \geq 1, \\
V(0,0) =& 0,\ V(0,S) = [1 - 1(1-p_s)^S]u_s,\ \forall S \geq 1,\ V(R,0) = \left[1-(1-p_r)^n\right]u_r, \forall R \geq 1.
\end{aligned}
\tag{9}
$$

Specifically, in the first equation of (9), the term $\frac{S}{R+S}\tilde{V}(R,S)$ corresponds to the case that the next arrival is a safety school (which happens with probability $\frac{S}{R+S}$), and the second term depicts the case wherein the next arrival is a reach school. In both scenarios, the outcomes can be either admission or rejection. The difference, as described before, is that a safety school may also reject the player, and the player would not necessarily decline it even if there are remaining safety schools. We can certainly develop a dynamic programming algorithm based on (9), and for conciseness we omit it here. Additionally, we can imitate the proof of Proposition 3 to show that $V(R,S)$ is discrete concave in R and S even when $p_s < 1$.

Proposition 8. *$\tilde{V}(n,k)$ is increasing in k, and is increasing and discrete concave in n, $\forall n$. Moreover, $p_s u_s \leq \tilde{V}(n,k) \leq u_r$, $\forall n$, $\forall k$. $V(R,S)$ is discrete concave in R and S.*

5.2 Numerical experiments: $p_s < 1$

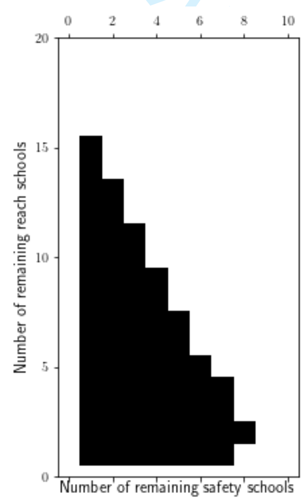

Figure 8: Range of paying the reservation fee. $u_r = 4.5, u_s = 3.0, p_r = 0.20, p_s = 0.5, F = 0.1$.

The complicated nature of the problem makes it difficult to fully solve the two-dimensional dynamic program in closed form. Thus, we exploit the discrete concavity and resort to numerical experiments to illustrate some key findings. We again focus on new features introduced in this framework. Specifically, we examine the critical decision making moment when the player faces a safety school's admission, namely the scenario for $\tilde{V}(n,k)$. The algorithm devel-

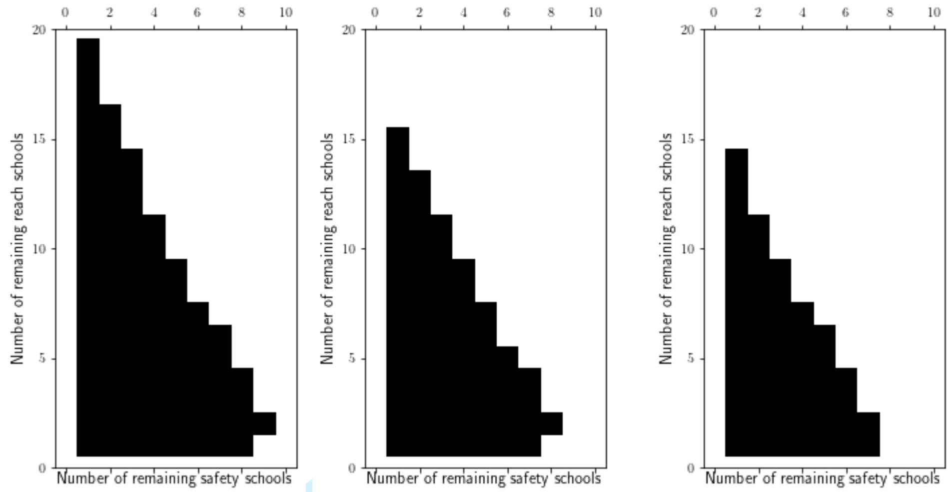

Figure 9: Range of paying the reservation fee. Left: $F = 0.05$; Middle: $F = 0.10$; Right: $F = 0.15$. Other parameter values are: $u_r = 4.5, u_s = 3.0, p_r = 0.2, p_s = 0.5$.

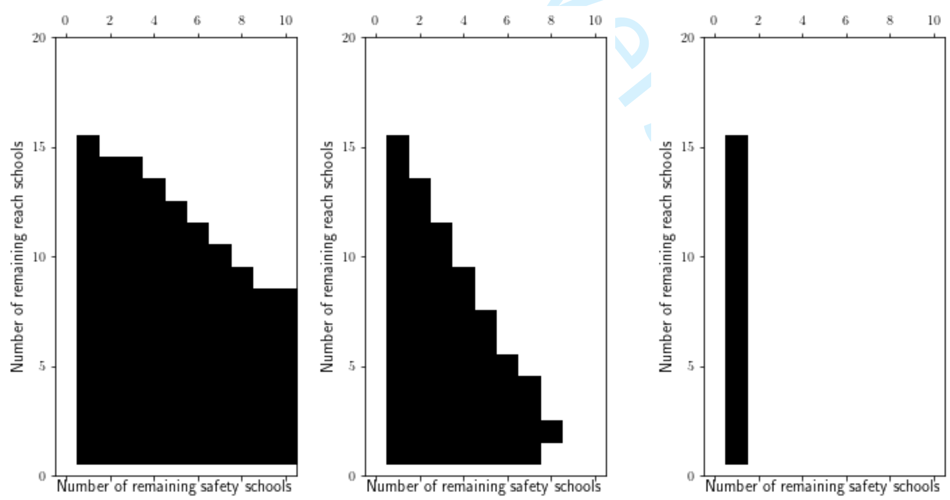

Figure 10: Range of paying the reservation fee. Left: $p_s = 0.20$; Middle: $p_s = 0.50$; Right: $p_s = 1.00$. Other parameter values are: $u_r = 4.5, u_s = 3.0, p_r = 0.2, F = 0.1$.

oped above is then used to calculate this, and we illustrate when the player chooses to pay the reservation fee F in Figures 8 –10. In Figure 8, we observe again that the range of paying the reservation fee can be non-monotone in S. In Figure 9, we vary the value of reservation fee F. We find that a larger F shrinks the range of paying it, and it can have differential impacts of different S. Figure 10 shows the range under different admission probabilities of safety schools p_s. The result suggests that a higher p_s reduces the range of paying F, and with deterministic safety school admissions the player waits for the last one.

In deriving the optimal choice among accepting, reserving, and declining above, we have already used the algorithm to obtain $\tilde{V}(n,k)$ and $V(R,S)$. We now show some patterns of optimal numbers of reach and safety schools, N_r^*, N_s^* in Figure 11. In all three examples, we vary the admission probability of safety schools p_s and see how the player optimizes her portfolio in stage 1. The three examples differ in the value of u_s, safety school payoff. The common feature is that the optimal number of reach schools decreases when it is more likely to get admissions from safety schools. This confirms our intuition as a higher p_s reduces the competitive advantage of reach schools. The subtlety again lies in the optimal number of safety schools: somewhat surprisingly the player may apply to fewer safety schools when the safety school admission is more secured. In the right panel of Figure 11, the optimal number of safety schools N_s^* can increase, decrease, increase, and finally decrease in probability p_s.

This non-trivial pattern arises from the joint forces of various factors. First, when the safety schools' admissions are more certain, these identical copies become more likely to be redundant since ultimately the player only needs to join one school (see the extreme case $p_s = 1$). Note that if there is no reach school, the incremental benefit of adding one more safety school is $(1-p_s)^{N_s} - (1-p_s)^{N_s+1}$. Second, a higher admission probability also makes the safety schools more appealing vis-a-vis reach schools. This opposite effect pushes the player to apply more safety schools. The third force is attributed to the interaction via uniform randomization. In Section 4.2 we have articulated the rationale of applying more safety schools: to postpone the decision-making stage such that the player has sufficient chances to learn reach schools' outcomes. When there are fewer reach schools due to increased p_s, this incentive drops as well; consequently, the player need not apply for so many safety schools to push the decision-making stage. While the second factor enhances the player's incentive to increase N_s^*, the first and third factors induce her to adjust it downwards.

6 Extensions

This section makes substantial relaxations of the key assumptions in our base model, and discusses the robustness and limitations of our analysis. We will concentrate on the general cases wherein $p_s < 1$ and elaborate on the first-stage (portfolio optimization) problems. A critical

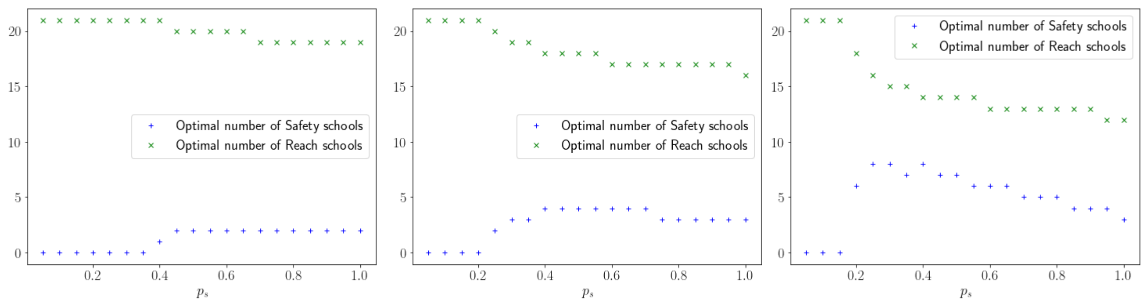

Figure 11: Optimal numbers of reach and safety schools. Left: $u_s = 2.5$; Middle: $u_s = 3.5$; Right: $u_s = 4.5$. Other parameter values are: $u_r = 5, p_r = 0.2, F = 0.1, a_r = 0.01, a_s = 0.004$.

step is to characterize the structural property of the player's stage 2 expected payoff $V(R, S)$, which excludes the application fees a_r, a_s. This notation $V(R, S)$ is adopted in all the following, even though we change the modeling details.

6.1 Correlated outcomes

While the independence assumption is adopted in most, if not all, of the published work on school application problems, this feature is not innocuous. Very recently, Ali and Shorrer (2025) show that in a simultaneous-search problem with correlated admission probabilities, the proof technique and the major findings in Chade and Smith (2006) no longer apply. Thus, it is sensible to incorporate the correlated admissions into our simultaneous-search sequential-outcome setting.

To wit, we will adopt the elegant framework proposed by Ali and Shorrer (2025). The micro-foundation arises from a common exam the player has taken, whose score will be used by schools to make admission decisions. Correlation among different schools' admission decisions arises from such a common exam, and each school has a different passing threshold (minimum score) for successful applications. Importantly, the player is uncertain about her own score at the time of applications, but will learn gradually the score through observing her application outcomes. Alternatively, the model can be interpreted as if the player knows her own score, but is uncertain about the admission criteria (minimum scores) set by those schools.

We consider two categories of schools: reach schools and safety schools, each of which has an infinite number of schools. Let $\tau_i \in [0, 1], i \in \{r, s\}$ represent the minimum score of a reach/safety school, and let ξ denote the player's score. Thus, the player's application to a reach school is accepted if and only if $\xi \geq \tau_r$, and likewise for a safety school's application $\xi \geq \tau_s$. In other words, $p_i = 1 - \tau_i, i \in \{r, s\}$, $p_r \equiv 1 - \tau_r < 1 - \tau_s \equiv p_s$, but the application outcomes are correlated through the common ξ. We maintain all other notation: payoffs u_r, u_s, application fees a_r, a_s, and reservation fee F.

With correlation outcomes, the new feature is that the player will update, at most twice, her belief about the schools yet to reveal their outcomes. Recall that $p_r < p_s$. Upon receiving a reach school's admission, all reach and safety schools will accept the player. A safety school's rejection implies that the player will be rejected by all reach and safety schools. The player's decision making is trivial in these two cases: accept the reach school's admission in the former and do nothing when seeing all rejections in the latter. Besides these two trivial cases, a reach school's rejection will lead to a downward belief adjustment: the player will also be rejected by all subsequent reach schools, and the belief for the subsequent safety schools' admission probability is $\frac{p_s - p_r}{1 - p_r}$. On the other hand, a safety school's admission will lead to an upward belief adjustment: the player will also be accepted by all subsequent safety schools, and the belief for the subsequent reach schools' admission probability is $\frac{p_r}{p_s}$. Since the uncertainty is resolved after learning from the first arrival of each type of school, the second stage problem becomes easier compared to the independent case in our base model. We now show in Proposition 9 that our approach can be adapted to handle this correlated scenario.

Proposition 9. *With correlation outcomes, $V(R, S)$ is discrete concave in R and S.*

6.2 Deposit

In this section, we consider the alternative setting: reservation and immediate acceptance both come along with deposit D, but this will be refunded to the player through tuition deduction once the player eventually joins the reserved school. In this setting, most of the results are the same as we modify how the reservation fee/deposit factors are included in the equations. We illustrate this via the toy example $(N_r, N_s) = (1, 1)$. Consider the case in which the safety school first and the reach school second. If the player keeps the safety school's admission, her utility is $p_r(u_r - D) + (1 - p_r)\,u_s$, where D is surrendered when the player joins the reach school instead. Back in period 1, the player needs to choose among the three options and obtain the maximum expected payoff among them:

$$\max\{u_s, -p_r D + p_r u_r + (1 - p_r)\,u_s, p_r u_r\}.$$

We now utilize the above properties for the general case $(N_r, N_s) = (R, S)$, and show that the procedure and dynamic programming structure remain valid.

Proposition 10. *With deposit D, $V(R, S)$ is discrete concave in R and S.*

Proposition 10 therefore implies that interested readers can efficiently characterize the optimal portfolio numerically using our algorithm.

6.3 Nonuniform randomization

Our uniform randomization assumption facilitates neat expressions of the probabilities for critical events, but the result can be generalized if certain conditions hold. To this end, we introduce the general function $G^R(N_r, N_s)$ as the probability of seeing a reach school as the next arrival when currently there are N_r reach schools and N_s safety schools yet to reveal their outcomes. Consequently, $G^S(N_r, N_s) \equiv 1 - G^R(N_r, N_s)$ is the probability that the next arrival is a safety school. In our base model with uniform randomization, $G^R(N_r, N_s) = \frac{N_r}{N_r+N_s}$, $\forall N_r, N_s$. As another example, if a reach school more likely reveals outcome earlier than a safety school, we can add weights λ_r, λ_s such that $G^R(N_r, N_s) = \frac{\lambda_r N_r}{\lambda_r N_r+\lambda_s N_s}$, $\forall N_r, N_s$, where $\lambda_r \geq \lambda_s > 0$.

With this modification, we note that Lemma 1 does not use these probability terms at all and readily applies. In the proofs we show the stochastic dominance of $\mathbb{P}(n)$ with respect to S via the nested probability events; the argument easily goes through beyond uniform randomization. Therefore, Propositions 1 and 2 hold true. When $p_s < 1$, the dynamic programming algorithm is valid too, so long as we modify the probability terms accordingly. The proofs of Propositions 3 and 8 rely on one feature: when we increase the number of reach schools, it is more likely that the next arrival is a reach school. We formally impose these assumptions in the following proposition such that the structural property of $V(R, S)$ remains valid.

Proposition 11. *With nonuniform randomization, suppose that $G^R(N_r, N_s)$ is increasing in N_r and $G^R(N_r, N_s)$ is discrete concave in $N_r, \forall N_r, N_s$. Then $V(R, S)$ is discrete concave in R and S.*

Proposition 11 shows that the approach applies well to nonuniform randomization as we modify those combinatorics formulas with the corresponding probabilities. The expressions are more involved, but some fundamental properties remain valid. As a remark, the conditions imposed on $G^R(N_r, N_s)$ ares automatically satisfied by uniform randomization in our base model and the generalized form $G^R(N_r, N_s) = \frac{\lambda_r N_r}{\lambda_r N_r+\lambda_s N_s}$.

6.4 Heterogeneous reservation fees

Another simplification is the common reservation fee, which can be relaxed as follows. Specifically, within each category we will keep the same admission probability, payoff, and application fee $p_i, u_i, a_i,\ i \in \{r, s\}$, but now let F_j denote the reservation fee of a safety school j. Per our discussion, the reservation fee of a reach school does not matter since this option is never exercised. Thus, reach schools with heterogeneous reservation fees are ex ante identical in the eyes of the player. On the other hand, there is an intuitive **pecking order** for the safety schools: if the player intends to apply for any safety school, she would rank them using the ascending order of reservation fees ($F_{(1)} \leq F_{(2)} \leq \ldots$), and go down the list. This can be immediately obtained via stochastic coupling arguments: replacing a safety school by another one with a strictly lower reservation fee can improve her payoff in every instance. As a side note, within the category of safety schools, increasing the reservation fee will unambiguously discourage applicants (cf. the finding from Figure 7).

If safety schools have heterogeneous reservation fees, it is no longer true that all safety schools except the last one are ignored. To illustrate, consider the following case. If some earlier safety schools have lower reservation fees or reservations are free, the player may reserve these admissions before moving to the last safety school. In such a scenario, we shall abandon the procedure in Sections 4.1 and 4.2 and use the first safety school approach in Section 5, or the dynamic programming recursion in (9). The iterative procedure to obtain $\tilde{V}(n,k)$ and $V(R,S)$ as described in our algorithm can be adjusted to handle this numerically.

Despite these issues, Proposition 12 shows that we can still maintain the well-behaved payoff function.

Proposition 12. *With heterogeneous reservation fees, $V(R,S)$ is discrete concave in R and S.*

Notably, the definitions of $\tilde{V}(n,k)$ and $V(R,S)$ have already incorporated the optimal decision rule of the aforementioned pecking order $F_{(j)}$ on safety schools, and this is crucial for establishing the discrete concavity in Proposition 12. This echoes our approach in Section 4.1: understanding the optimal decision rule helps us tackle the structural properties of dynamic programs.

7 Conclusions and discussions

In this paper, we use postgraduate program applications to illustrate the key ingredients of a family of problems: simultaneous search with sequential outcomes, and we provide rigorous analysis for the first-of-its-kind settings. A single player faces two categories of schools: reach and safety schools. The decision making is done in two stages. In the first stage, the player first decides the application portfolio. In the second stage, the outcomes from the schools applied to arrive randomly over time, and the player must pay reservation fees to maintain the eligibility for recalling the earlier offers. When the safety schools always admit the player, it suffices to focus on the last safety school. The player's payoff after applications is increasing and discrete concave in the number of safety schools, and the optimal number of reach schools is increasing in the reservation fee. When the admissions of safety schools are no longer guaranteed, we find that once the player pays the reservation fee for one safety school's admission, she will ignore all the subsequent safety schools. These hint to a simple dynamic programming algorithm to solve the problem in its entirety. Out analysis extends to a few variants of model setups, including correlated outcomes among schools, deposits rather than reservation fees, nonuniform randomization of outcome revelations, and heterogeneous reservation fees.

One might be tempted to conjecture that when either the reservation fee increases or the admission probability is low, the player shall apply to fewer safety schools. However, we utilize our dynamic programming algorithm to demonstrate numerically that this is not always true. We can find counterexamples in which (1) the player applies to more safety schools when the reservation fee gets higher or the probability of admission to safety schools is lower; and (2) the optimal number of safety schools is nonmonotone in either the reservation fee or the admission probability. This has strong managerial implications for service providers in devising their reservation fees and deposits, especially for those institutions that are not unambiguously chosen by prospective applicants.

Our efforts serve as an initial attempt to tackle this complicated problem and naturally many possible extensions are in order. For example, when an individual applicant faces such a problem, one have to consider that other applicants may experience similar situations at the same time. Since these applicants compete for limited spots in the postgraduate programs, their decisions may interfere and such externality is absent in our simplified model. Another possibility is that safety schools may anticipate how applicants make decisions and proactively adjust the appropriate timing for their outcome revelations. This school-side deliberation is also left out in the current analysis.

While we motivate this research via postgraduate programs, our framework has wide applications. As indicated by Chade and Smith (2006), the stage-one simultaneous-search problem is appropriate when a firm chooses among available technologies and has to race against competitors for timely explorations. As another example, an institution may invite job candidates for campus visits and rank candidates accordingly (Galenianos and Kircher (2009), Hu and Tang (2021), and Zorc and Tsetlin (2020)). In the online labor markets such as Amazon Mechanical Turk, requesters can post jobs that are available to a group of eligible workers; targeting a specific worker and moving through candidates one by one is not feasible given the current platform setup. De los Santos et al. (2012) empirically validate that when consumers purchase products either online or offline, their search behaviors can be suitably described by the simultaneous-search setup. Thus, our analysis applies to all these contexts beyond the postgraduate programs.

Our framework introduces the novel aspect through stage two that there is no synchronized outcome revelation stage among available choices, and the decision maker has no discretion in pushing the applied entities to make decisions immediately. This would occur, for example, when the firm explores technologies, these R&D teams or start-up tech companies may reply with positive or negative news, and the responses are not synchronized. The firm may have to make early-round investments to start-up companies to secure the future opportunity of acquisition and still keep an eye on other start-ups. The phenomenon also arises when the committee decisions from different institutions occur sequentially, or when consumers consult their friends about product user experiences and these friends respond sporadically. In Amazon Mechanical Turk, a priori requesters do not know which workers will complete the tasks for payments, and depending on the responses, requesters may determine when to close the online labor recruitment. Our framework allows for a full spectrum of reservation fees, thereby accommodating both perfect recall and no recall as special cases.

Acknowledgement

I thank Michael Pinedo (the department editor) for the continuing guidance and support throughout the project's development, and the senior editor and the reviewers for their detailed comments and many valuable suggestions that have significantly improved the quality of the paper. I have also benefited from the conversation with Yufeng Guan, Zhenyu Hu, Wenjuan Li, Zhixi Wan, Qiuyi Yan, and Xinyi Zhou, and conference and seminar participants in Analytics for X Conference, Frontiers in Operations Management Workshop, Next-Gen Scholars Symposium, CityU, CUHK, HKU, HKUST, Lingnan, NUS, and SUSTech. This project is partially funded by Hong Kong RGC C6020-21GF and Crown Worldwide Professorship. All remaining errors are my own.

References

Ali, S. N. and R. I. Shorrer (2025). Hedging when applying: Simultaneous search with correlation. *American Economic Review 115*(2), 571–598.

Anderson, E., B. Chen, and L. Shao (2017). Supplier competition with option contracts for discrete blocks of capacity. *Operations Research 65*(4), 952–967.

Armstrong, M. (2017). Ordered consumer search. *Journal of the European Economic Association 15*(5), 989–1024.

Barthel, A.-C. and T. Sabarwal (2018). Directional monotone comparative statics. *Economic Theory 66*(3), 557–591.

Boyd, S. and L. Vandenberghe (2004). *Convex Optimization*. New York: Cambridge University Press.

Burdett, K. and T. Vishwanath (1988). Declining reservation wages and learning. *The Review of Economic Studies 55*(4), 655–665.

Carrasco, J. A. and L. Smith (2017). Search at the margin. *American Economic Review 107*(10), 3146–3181.

Chade, H., J. Eeckhout, and L. Smith (2017). Sorting through search and matching models in economics. *Journal of Economic Literature 55*(2), 493–544.

Chade, H. and L. Smith (2006). Simultaneous search. *Econometrica 74*(5), 1293–1307.

Chen, X. and M. Li (2021a). Discrete convex analysis and its applications in operations: A survey. *Production and Operations Management 30*(6), 1904–1926.

Chen, X. and M. Li (2021b). M-convexity and its applications in operations. *Operations Research 69*(5), 1396–1408.

Choi, M., A. Y. Dai, and K. Kim (2018). Consumer search and price competition. *Econometrica 86*(4), 1257–1281.

Daughety, A. F., J. F. Reinganum, et al. (1992). Search equilibrium with endogenous recall. *RAND Journal of Economics 23*(2), 184–202.

De los Santos, B., A. Hortaçsu, and M. R. Wildenbeest (2012). Testing models of consumer search using data on web browsing and purchasing behavior. *The American Economic Review 102*(6), 2955–2980.

Doval, L. (2018). Whether or not to open Pandora's box. *Journal of Economic Theory 175*(C), 127–158.

Eeckhout, J. and P. Kircher (2010). Sorting and decentralized price competition. *Econometrica 78*(2), 539–574.

Galenianos, M. and P. Kircher (2009). Directed search with multiple job applications. *Journal of Economic Theory 144*(2), 445–471.

Hu, Z. and W. Tang (2021). Size matters, so does duration: The interplay between offer size and offer deadline. *Management Science 67*(8), 4935–4960.

Hu, Z. and Y. Xiao (2024). Optimal dynamic mechanism under customer search. Operations Research. Forthcoming.

Liu, Y., W. L. Cooper, and Z. Wang (2019). Information provision and pricing in the presence of consumer search costs. *Production and Operations Management 28*(7), 1603–1620.

Petrikaitė, V. (2018). Consumer obfuscation by a multiproduct firm. *The RAND Journal of Economics 49*(1), 206–223.

Rogerson, R., R. Shimer, and R. Wright (2005). Search-theoretic models of the labor market: A survey. *Journal of Economic Literature 43*(4), 959–988.

Smith, L. (2006). The marriage model with search frictions. *Journal of Political Economy 114*(6), 1124–1144.

Stigler, G. J. (1961). The economics of information. *Journal of Political Economy 69*(3), 213–225.

Weitzman, M. L. (1979). Optimal search for the best alternative. *Econometrica 47*(3), 641–654.

Wu, C., C. Jin, and Y.-J. Chen (2022). Managing customer search via bundling. *Manufacturing & Service Operations Management 24*(4), 1906–1925.

Zorc, S. and I. Tsetlin (2020). Deadlines, offer timing, and the search for alternatives. *Operations Research 68*(3), 927–948.