
ECONtribute

Discussion Paper No. 434

Pandora's Box Reopened: Robust Search and Choice Overload

Sarah Auster

Yeon-Koo Che

October 2026

www.econtribute.de



**UNIVERSITÄT
ZU KÖLN**

Pandora’s Box Reopened: Robust Search and Choice Overload*

Sarah Auster Yeon-Koo Che[†]

August 18, 2026

Abstract

This paper revisits the classic Pandora’s box problem, studying a decision-maker (DM) who seeks to minimize her maximal ex-post regret. The DM decides how many options to explore and in what order, before choosing one or taking an outside option. We characterize the regret-minimizing search rule and show that the likelihood of opting out often increases as more options become available for exploration. We show that this “choice overload” is driven by the DM’s fear of *selection error*—the regret from searching the wrong options—suggesting that steering choice via recommendations or cost heterogeneity can mitigate regret and encourage search.

Keywords: Sequential Search, Minmax Regret, Choice Overload

1 Introduction

The emergence of digital markets has revolutionized the way individuals make choices, offering unprecedented access to a vast array of products, services, and information. While this abundance empowers consumers, it also presents a significant challenge: navigating a

*We thank Pak Hung Au, Ben Brooks, Laura Doval, Yingni Guo, Emir Kamenica, Teddy Kim, Andreas Kleiner, Suraj Malladi, Pietro Ortoleva, Ludvig Sinander, Bruno Strulovici, Leeat Yariv, Andy Zapechelnnyuk, and audiences in various seminars and conferences for helpful comments and discussions. Sarah Auster gratefully acknowledges funding from the European Research Council (ERC)—grant No 101165999, the German Research Foundation (DFG) through Germany’s Excellence Strategy—EXC 2126/1—390838866 and CRC TR 224 (Project B02). Yeon-Koo Che is supported by the Ministry of Education of the Republic of Korea and the National Research Foundation of Korea (NRF-2024S1A5A2A0303850912)

[†]Auster: Department of Economics, University of Bonn (email: auster.sarah@gmail.com); Che: Department of Economics, Columbia University (email: yeonkooche@gmail.com)

complex choice environment where the value of alternatives is not immediately known and can only be discovered through costly effort. This challenge is starkly evidenced by the phenomenon of “choice overload,” where an overabundance of options can paradoxically discourage active selection and diminish satisfaction (see [Scheibehenne et al. \(2010\)](#) and [Chernev et al. \(2015\)](#) for reviews).¹

The canonical framework for analyzing such problems of costly discovery is the sequential search model pioneered by [Weitzman \(1979\)](#). This model provides an elegant and simple index-based rule for an expected utility-maximizing decision-maker (DM) who possesses a clear, *unambiguous* prior about the *independent* value distribution of different alternatives. In many real-world settings, however, individuals lack clear probabilistic perceptions of the merits of choices, particularly when they are complex and unfamiliar. Instead, they face ambiguity, also known as Knightian uncertainty. We therefore depart from the standard assumption by considering an individual facing such uncertainty.

Specifically, we adapt Weitzman’s classic Pandora’s box problem, where a DM faces n alternatives, or “boxes,” with unknown rewards. For most of the analysis, we focus on the binary case in which boxes either contain a high reward $\bar{u}$ or nothing. Opening box i and observing its reward costs c_i , with boxes indexed in increasing order of costs, $c_1 < \dots < c_n$. The DM may inspect boxes sequentially and, after each inspection, decides whether to continue searching or to stop. Upon stopping, she chooses the best alternative found so far, or takes an outside option with payoff zero. Unlike in the standard Weitzman model, the DM does not know the joint distribution of rewards, $F \in \Delta(\{0, \bar{u}\}^n)$, but considers all such distributions possible, allowing arbitrary correlation across the rewards of different boxes. Restrictions such as i.i.d. rewards are interpreted as additional information and considered separately below.

We assume that the DM confronts this ambiguity by seeking to minimize her **maximal anticipated ex-post regret**. The minmax regret criterion is a compelling objective for individuals concerned with foregone opportunities.² Rather than reflecting an emotional bias, this approach can be understood as a robust decision rule for an agent who cannot or will not adopt a single prior, effectively guiding her to hedge against the scenarios where her choices are most consequential. Analytically, the problem amounts to a zero-sum game

¹This counterintuitive behavior has been documented in classic field experiments, such as shoppers being less likely to purchase jam from a selection of 24 varieties than from a selection of 6, and in high-stakes decisions, where employees are less likely to enroll in retirement plans with more fund options ([Iyengar and Lepper, 2000](#); [Sethi-Iyengar et al., 2004](#)).

²Another robust approach is the maxmin criterion, under which the DM maximizes her payoff against the worst-case distribution of rewards. In the present setting, however, this solution concept does not capture a meaningful tradeoff: given the degenerate worst-case distribution of zero rewards, it yields a trivial “universal no-search” policy.

between the DM and an adversarial “Nature.” For any search strategy the DM chooses, Nature responds by selecting the reward distribution that maximizes the DM’s expected regret. The DM’s regret is measured relative to an oracle—an ideal self who maximizes the payoff after observing realized rewards in all boxes. The DM’s optimal search strategy is the one that minimizes the maximal regret.

Our first result shows that Nature’s worst-case scenario takes a stark and tractable form. For any search rule chosen by the DM, there is a regret-maximizing distribution that places the high reward in *at most* one box. The robust search problem is therefore a “needle-in-the-haystack” problem. Nature need not make the world rich in opportunities; it is enough to make the DM worry that either there is no prize at all or a single prize is hidden in the box she is least likely to inspect. This reduction is important both economically and technically. Economically, it identifies the source of regret as a selection problem: with many boxes, the DM is not only concerned about searching too much or too little, but also about searching in the wrong place. Technically, it replaces the original ambiguity over all binary reward vectors by a game with $n + 1$ states.

Building on this insight, we characterize the DM’s robustly optimal search rule. We show that the rule partitions the choice set into *contiguous blocks* of boxes with similar search costs. The DM searches these blocks sequentially in order of increasing cost and treats each block as an all-or-nothing unit: upon reaching a block, she either stops or searches the entire block, randomizing over the order of boxes within any non-degenerate block. This structure equalizes regret across all boxes under consideration. Although cheaper boxes generate greater regret when the DM stops before reaching them, they are also searched earlier in expectation and therefore entail a lower expected search cost. When the cost difference between two adjacent boxes is sufficiently large, regret equalization requires the DM to open the cheaper box before the more expensive one with certainty, placing the two boxes in separate blocks. By contrast, when costs are similar, a deterministic ordering would impose excessive expected search costs when the prize is located in the later box. In that case, the boxes are grouped into the same block, and their search order is randomized.

Assuming that cost differences become smaller as costs increase, the general partition structure takes the form of a simple tail-bunching rule. Low-cost boxes are separated and searched sequentially, with random stopping between them, while high-cost boxes are bunched into a single tail block and searched exhaustively if the tail is reached. Behaviorally, the DM hesitates to continue search after each inspected option when the menu is large, but searches all remaining options once the set has been reduced to a manageable size.

The main technical challenge in the analysis is to characterize when regret equalization within a block can be implemented by a stochastic search order, formally a lottery over

permutations of that block. After deriving a target vector of expected cumulative search costs for each location in a given block, we show that this vector is feasible if and only if no subset of boxes is assigned a lower weighted cumulative search cost than it could obtain by being searched first. The necessity of these conditions is immediate. Sufficiency follows by interpreting the set of allocations satisfying these inequalities as the core of a hypothetical cooperative game. We show that this game is supermodular and apply [Edmonds \(2003\)](#); [Ichiishi \(1981\)](#) to establish that every element of the core is implementable by some randomization over search orders.

The comparative statics link the model back to the challenges consumers face when navigating rich choice sets in digital markets. We show that adding new options always reduces the DM’s willingness to initiate search, provided that the search costs of the new options are at least as high as that of the cheapest existing option. The model thus predicts a higher likelihood of opting out as the choice set expands, a key marker of choice overload. Holding fixed the true data-generating process, we also show that the DM’s reduced inclination to initiate search generates the familiar hump-shaped relationship between menu size and the probability that the DM ultimately chooses an option from the menu (e.g., purchases a product from some seller): this probability increases for low values of n , but falls for large values of n . Intuitively, each additional option creates a new possible location for the best product and therefore expands the set of ways in which the DM may make a selection error. While such hump-shaped purchase patterns are observed both in experimental and online retail settings ([Shah and Wolford, 2007](#); [Long et al., 2025](#)), they are difficult to reconcile with Bayesian decision-making, where added options can be ignored and therefore weakly increase expected search.

The problem of not knowing where to start is exacerbated when boxes become more homogeneous ex ante, that is, when their search costs are similar. While a steep cost profile provides some guidance about where to begin, a flat profile leaves the DM facing a more symmetric and therefore more ambiguous navigation problem. The homogeneous limit makes this force stark. We show that, as costs converge to a common constant $c \in (0, \bar{u})$, one of two possibilities arises. If n is sufficiently small, all boxes are grouped into a single block, which the DM searches exhaustively with positive probability. If, instead, n is sufficiently large, the probability of initiating search converges to zero as the cost profile becomes homogeneous. Importantly, in this second regime, the DM’s limiting worst-case belief assigns probability one to the prize being in one of the boxes. Hence, the DM withdraws not because she thinks there is nothing to find, but because she lacks guidance about how to find it. This provides a clear interpretation of the digital-market paradox: online environments lower the cost of inspecting any single item, but they also create large, weakly differentiated menus in which

the main burden is deciding how to navigate them.

Finally, we show that the main mechanism is robust to several extensions. We first relax the binary-reward assumption and allow rewards to take any value in the bounded interval $[0, \bar{u}]$. While the analysis becomes more complex, the core overload result continues to hold: sufficiently large homogeneous menus make immediate opt-out robustly optimal. We next show that, in the homogeneous-cost case, our results are robust when we allow dynamic reoptimization, thereby relaxing commitment, and when we add the restriction that rewards are *independently* and *identically* distributed across boxes.

From an applied perspective, our work advances the understanding of how individuals navigate choice-rich environments with non-quantifiable uncertainty. The model’s predictions align with empirical regularities in sequential search that are difficult to reconcile with the standard model—most notably under-searching (e.g., Schotter and Braunstein, 1981; Cox and Oaxaca, 1989; Sonnemans, 1998; Kogut, 1990; Einav, 2005) and menu-size effects (de Lara and Dean, 2024). Regret in our setting can arise along two margins: an *intensity margin* (“how many options to inspect”) and a *selection margin* (“in which order to inspect them”). Our results show that the selection margin plays a central role: the refusal to search can be traced to the DM’s uncertainty about the “right” search order—an effect that is particularly strong when the options are ex ante homogeneous and offer little guidance on where to begin. This diagnosis yields concrete implications for *assortment design*: platforms can mitigate overload by deliberately breaking symmetry, for example by creating cost heterogeneity or using recommendations to guide the search order.

Beyond these design implications, the analysis is interesting from a methodological point of view. While the optimal search rule in settings with independently distributed rewards is well understood, results for more general environments, allowing for arbitrary correlations across values of different choice options, have remained elusive. Our approach, based on worst-case regret minimization, proves highly tractable, paving the way for the study of new economic questions in sequential search.

Related Literature. Our work is situated at the intersection of the literatures on sequential search, decision-making under ambiguity, and the behavioral phenomenon of choice overload. The analysis directly revisits the canonical Pandora’s box problem of Weitzman (1979), which provides the foundational model of optimal search for an expected-utility maximizer with known, independent priors. While this framework has generated a large and active literature in economics (e.g., Olszewski and Weber (2015); Doval (2018); Choi and Smith (2020)), existing studies typically adopt a Bayesian approach, in contrast to the present paper. Meanwhile, computer science and operations research on the subject focus

on the regret guarantees of given algorithms in rich settings rather than the optimal one in a more stylized environment like ours. Examples include [Beyhaghi and Kleinberg \(2019\)](#), [Chawla et al. \(2020\)](#), [Chawla et al. \(2021\)](#), or [Gergatsouli and Tzamos \(2024\)](#) (see [Beyhaghi and Cai \(2024\)](#) for a recent overview).

Our primary departure is to introduce ambiguity, connecting our paper to other recent work that relaxes Weitzman’s assumptions. Closest to our work, [Schlag and Zapechelnyuk \(2021\)](#) study minmax regret in an *infinite-horizon* sequential search problem. Their environment is not a finite-menu Pandora problem with option-specific inspection costs: search is undirected, and time is modeled through discounting rather than heterogeneous inspection costs. In a setting of *spatial search*, [Malladi \(2025\)](#) studies a regret-minimizing DM who learns about rewards, knowing that nearby options have similar quality. [Malladi \(2025\)](#) shows that relaxing the restriction on how much rewards can vary across locations—thereby increasing the complexity of the problem—can eliminate search incentives, mirroring our search-breakdown result when rewards are arbitrarily correlated.³

While our contribution is to the theory of search under ambiguity, our model provides a new theoretical foundation for the well-documented *choice overload* phenomenon ([Iyengar and Lepper, 2000](#); [Sethi-Iyengar et al., 2004](#); [Dean et al., 2022](#)).⁴ Other theoretical explanations for this behavior have focused on consumer inference and steering under Bayesian choice ([Kamenica, 2008](#); [Nocke and Rey, 2024](#)) or reference dependence ([Deb and Zhou, 2018](#)). Our paper joins the stream of literature that links choice overload to anticipated regret, for which decision theorists have provided axiomatic foundations ([Sarver, 2008](#); [Burturak and Evren, 2017](#)). We contribute a tractable search model with a standard, symmetric regret formulation that generates choice overload as a direct consequence of ambiguity about the menu of options. By contrast, [de Lara and Dean \(2024\)](#) study a rational expected-utility (EU) maximizer who is uncertain about the parameters of the reward distribution. This leads to a model of Bayesian learning and an effective positive correlation between options, whereas our regret-minimizing agent is guided by a worst-case belief that induces a negative correlation, generating distinct behavioral dynamics.

³In his setting, the worst-case scenario is a quality distribution over the exploration space with a “high” peak concentrated on a small interval of locations, which can be viewed as a spatial analogue of our “needle-in-a-haystack” structure. See also [Banchio and Malladi \(2025\)](#).

⁴Other examples include [Dean et al. \(2008\)](#), [Haynes \(2009\)](#), [Iyengar and Kamenica \(2010\)](#), and [Inbar et al. \(2011\)](#).

2 Setting

We consider the following Pandora’s box problem. A DM can open up to n boxes before choosing one of them or taking an outside option with payoff zero (perfect recall). Boxes are indexed by $i \in [n] := \{1, \dots, n\}$ and deliver a reward $u_i \in \mathbb{R}$. We write a state as $\omega = (u_1, \dots, u_n)$. For most of the analysis, we will focus on a binary reward setting where $u_i \in \{0, \bar{u}\}$ with $\bar{u} > 0$ —an option is a match or no match. The state space is thus

$$\Omega = \{0, \bar{u}\}^n.$$

The restriction to binary rewards gives us sufficient tractability to characterize the structure of robust search. As we will see in [Section 5.1](#), the main behavioral forces extend to the case where rewards are only restricted by an upper bound and, thus, may be continuously distributed.

To learn the reward of box i , the DM has to pay a search cost $c_i \in (0, \bar{u})$. We assume boxes have different inspection costs and index them in increasing order of their search costs:

$$i < j \iff c_i < c_j.$$

The DM decides sequentially whether to open one remaining box or stop. If no high-reward box has been discovered, stopping yields payoff zero. If a high-reward box has been discovered, it is optimal to stop, since no remaining box can deliver a payoff above $\bar{u}$ and search costs are positive.

EU benchmark. [Weitzman \(1979\)](#) considered the case of independent rewards with a known distribution, showing that the optimization problem for a standard EU maximizer has a cutoff solution: assigning to each box a reservation value, representing the minimum value that justifies the opening cost, the optimal strategy is to open boxes in descending order of this value until the best reward found so far exceeds the reservation values of all unopened boxes. In the case of binary rewards, if p_i denotes the probability that box i delivers reward $\bar{u}$, box i ’s reservation value is

$$z_i = \bar{u} - \frac{c_i}{p_i}.$$

Thus, before a high reward has been discovered, the DM opens boxes in descending order of $\bar{u} - c_i/p_i$ as long as it is positive for some remaining box.

While the case of independently distributed rewards has a simple and well-understood

solution, little is known about this fundamental problem in the general case allowing correlation.⁵

Robust search. Our key assumption is that the DM does not know how rewards are distributed. Rather than maximizing expected utility against a known prior, she evaluates a search rule by its maximal ex-post regret. The DM considers all distributions over Ω possible, so the baseline ambiguity set is $\Delta(\Omega)$. This allows for arbitrary correlation across boxes and imposes no probabilistic model on the reward-generating process.

For any search rule a , the benchmark for regret is an *oracle* who knows the realized state $\omega \in \Omega$. Let $u(a, \omega)$ denote the DM’s expected payoff from search rule a (potentially stochastic) in state ω , net of search costs, and let

$$a^*(\omega) \in \arg \max_{a'} u(a', \omega)$$

denote an oracle strategy. Since $c_i < \bar{u}$ for every i , the oracle payoff is given by

$$u(a^*(\omega), \omega) = \begin{cases} 0, & \text{if } u_i = 0 \text{ for all } i \in [n], \\ \bar{u} - \min\{c_i : u_i = \bar{u}\}, & \text{otherwise.} \end{cases}$$

The regret of strategy a in state ω is therefore

$$u(a^*(\omega), \omega) - u(a, \omega).$$

Strategies and payoffs. In the case of binary rewards, it is optimal for the DM to stop the search when discovering a high-reward box. Hence, it is without loss to describe the DM’s strategy conditional on not having found the high reward so far. Let $\mathcal{N} := 2^{[n]}$ denote the set of all subsets of boxes. A pure contingent search plan is a mapping

$$A = \left\{ \bar{a} : \mathcal{N} \rightarrow [n] \cup \{0\} \mid \bar{a}(N) \in N \cup \{0\}, \forall N \subseteq [n] \right\},$$

where $\bar{a}(N) = i \in N$ means that the DM opens box i when the set of remaining boxes is N , and $\bar{a}(N) = 0$ means that she stops. A mixed search rule is an element of $\Delta(A)$.

The DM’s goal is to minimize the maximum regret over all states in Ω . A **robustly**

⁵Weitzman’s index policy is a special case of the Gittins index policy. With correlated values, observing one box changes beliefs about the remaining boxes, so—absent additional structure—the simple state-independent index characterization is generally lost (see [Rothschild \(1974\)](#)); economics treatments of such non-independent environments include [Keller and Oldale \(2003\)](#); [Ke and Lin \(2020\)](#); [Bao et al. \(2023\)](#); [Gambato \(2024\)](#); [Bizzarri and Lomys \(2024\)](#).

optimal search rule thus solves the following problem:⁶

$$\min_{a \in \Delta(A)} \max_{\omega \in \Omega} \left[u(a^*(\omega), \omega) - u(a, \omega) \right]. \quad (1)$$

The solution to (1) is a Nash equilibrium—a *saddle point*—of a zero-sum game between the DM, who chooses a search rule to minimize regret, and an adversarial “Nature”, who chooses a reward assignment to maximize regret.⁷ This saddle point is typically in mixed strategies. The set of mixed strategies for Nature is the simplex $\Delta(\Omega)$, i.e., the set of all reward distributions including those featuring arbitrary correlations. Defining $\mathcal{R} : \Delta(A) \times \Delta(\Omega) \rightarrow \mathbb{R}$ as the regret value for fixed strategies of the players in this hypothetical game, the pair (a^*, F^*) is a saddle point if

$$\mathcal{R}(a, F^*) \geq \mathcal{R}(a^*, F^*) \geq \mathcal{R}(a^*, F), \quad \text{for all } a \in \Delta(A), F \in \Delta(\Omega).$$

The first inequality implies that, against the saddle-point belief F^* , the DM’s saddle-point strategy a^* is optimal in the standard sense: when fixing Nature’s choice, the expected oracle payoff is fixed, so minimizing expected regret is equivalent to maximizing expected utility.

3 Analysis and Main Results

We now turn to the solution of the robust search problem. The analysis begins with the one-box case, which isolates the basic minmax-regret tradeoff. We then move to the multiple-box case, where a new force appears: the DM must decide not only how much to search, but also where to search and in what order.

Single box $n = 1$: We start with the simplest version of the problem, with a single box. Let the cost of opening the box be c , and let $\alpha_1 \in [0, 1]$ denote the probability with which the DM opens it. In the saddle-point representation, write p for Nature’s probability that the box contains the high reward.

There are two ways in which the DM can regret her decision. If the box contains no prize, the oracle takes the outside option and obtains zero. The DM regrets opening the box, because opening only wastes the search cost. This regret is

$$R(\alpha_1; p = 0) = \alpha_1 c.$$

⁶This formulation of the problem assumes that the DM can commit to a dynamic search rule, which is not without loss under the minmax regret criterion. We return to this issue in [Section 5.2](#).

⁷Existence follows from the minmax theorem for finite zero-sum games: A and $\Omega = \{0, \bar{u}\}^n$ are finite, and the regret payoff is bilinear in mixed strategies.




Figure 1: The one-box problem. The increasing line is regret when the box contains no prize; the decreasing line is regret when the box contains the prize. The opening probability α_1^* that equalizes these two forms of regret minimizes the maximal regret, depicted by the green upper envelope.

If instead the box contains the prize, the oracle opens it and obtains $\bar{u} - c$. The DM now regrets not opening the box. This regret is

$$R(\alpha_1; p = 1) = (1 - \alpha_1)(\bar{u} - c).$$

Thus, increasing the search probability reduces regret in the prize state but increases regret in the no-prize state. Against a fixed p , expected regret is

$$R(\alpha_1; p) = (1 - p)\alpha_1 c + p(1 - \alpha_1)(\bar{u} - c).$$

Nature's pure states correspond to the two extreme lines in **Figure 1**: $p = 0$ generates the regret from searching in vain, while $p = 1$ generates the regret from missing the prize. The DM's minmax choice equalizes these two regrets:

$$\alpha_1^* = \frac{\bar{u} - c}{\bar{u}}.$$

Intuitively, the DM randomizes in order to hedge between two possible ex-post mistakes: searching when there is nothing to find and failing to search when the prize is present. The corresponding worst-case regret and Nature's saddle-point probability assigned to the prize

state are

$$R_1^* = \frac{c(\bar{u} - c)}{\bar{u}}, \quad \hat{p} = \frac{c}{\bar{u}}.$$

At the probability $\hat{p}$, the DM is indifferent between opening and not opening: opening risks paying c in the no-prize state, while not opening risks losing $\bar{u} - c$ in the prize state. In the figure, the line for $p = \hat{p}$ is flat, and all expected-regret lines pass through the same minmax point.

Multiple boxes $n \geq 2$: With several boxes, the previous tradeoff remains, but a new margin appears. The DM must decide not only whether to search, but also in what order to search. A wrong order creates regret even when the DM eventually finds the prize. If the prize is in a cheap box but the DM first opens an expensive or unproductive box, the oracle would have gone directly to the cheap prize box and saved the extra cost. Thus the multi-box problem combines the one-box stopping tradeoff with a selection problem: the DM must guard against searching too much, searching too little, and searching the wrong boxes.

From a technical point of view, the key complication with multiple boxes is that rewards may be arbitrarily correlated, rendering the standard Weitzman index rule inapplicable and the resulting problem seemingly intractable. Crucially, imposing the robustness criterion restores tractability: our first result establishes that, despite the large underlying state space, Nature’s worst-case distribution takes a particularly simple form.

Lemma 1 (Needle in the haystack). *Fix any search rule a . Then there is a regret-maximizing distribution that assigns $\bar{u}$ to at most one box.*

The intuition is simple. Consider any state with several high-reward boxes, and let $i^*(\omega)$ be the cheapest box among those containing the high reward. Now remove all other high rewards, leaving only the prize in box $i^*(\omega)$. The oracle’s payoff is unchanged, because the oracle would open the cheapest prize box in either state. The DM’s payoff, however, can only fall. Removing prizes cannot make the DM find a prize earlier, and may make her search longer or fail to find a prize at all. Hence regret weakly increases after all but the cheapest prize are removed. Nature can therefore restrict attention to states with no prize or exactly one prize.

As a consequence, we can replace the original state space $\{0, \bar{u}\}^n$ by the reduced state space

$$\{0, 1, \dots, n\},$$

where state 0 means that no box contains a prize and state $i \in [n]$ means that box i is the

unique prize box. Nature's mixed strategy is then denoted by

$$\mathbf{p} = (p_0, p_1, \dots, p_n) \in \Delta(\{0, 1, \dots, n\}).$$

In reduced state $i \geq 1$, the oracle opens box i immediately, so its payoff is $\bar{u} - c_i$. The oracle's expected payoff under $\mathbf{p}$ is therefore

$$\sum_{i=1}^n p_i (\bar{u} - c_i).$$

The one-prize reduction is economically important. It says that the worst case is not a world with many hidden opportunities. It is a needle-in-the-haystack environment.⁸ Nature makes the DM worry that there may be no prize at all, and, if there is one, that it may be hidden in the box the DM is least inclined to search. The rest of the analysis characterizes how the DM optimally hedges against this possibility.

3.1 The Block Structure of the Saddle Point

We now establish the main structure of the saddle points of the zero-sum game. We refer to *active boxes* as the boxes to which Nature, under her saddle-point strategy, assigns the prize with positive probability.

We say that a rule searches a block of boxes *exhaustively* if, conditional on reaching the block without having found a prize, it continues until either a prize is found or every box in the block has been opened.

Theorem 1. *In any saddle point, the active boxes form an initial segment $[m] \subseteq [n]$, which can be partitioned into maximal contiguous blocks*

$$B^{(1)}, B^{(2)}, \dots, B^{(K)}$$

such that

$$\frac{p_i}{c_i} = \lambda^{(k)} \quad \text{for all } i \in B^{(k)}, \quad \lambda^{(1)} > \lambda^{(2)} > \dots > \lambda^{(K)}.$$

Moreover, every pure rule in the support of the DM's mixed strategy searches blocks in the order

$$B^{(1)}, B^{(2)}, \dots, B^{(K)},$$

and, for each block, either does not enter the block or searches it exhaustively.

⁸The same prize structure is studied in [Kamenica \(2008\)](#); in his setting, however, it is assumed exogenously rather than arising endogenously as the worst-case scenario, as it does in ours.

The theorem says that robust search has a simple architecture. Boxes are ordered according to their costs and then grouped into blocks. If the active set is $[m]$, then there are endpoints $0 = n_0 < n_1 < \dots < n_K = m$ such that

$$B^{(k)} = \{n_{k-1} + 1, \dots, n_k\}, \quad k = 1, \dots, K.$$

The DM searches block by block, possibly exiting between blocks and randomizing the order within non-degenerate blocks. Under a saddle-point strategy, any active block is searched with positive probability and contains the prize with positive probability. High-cost boxes outside $[m]$, if any, are ignored by both the DM and Nature.

Outline of the proof. The proof of [Theorem 1](#) has two main steps. In the first step, we characterize the DM's best reply to an arbitrary one-prize structure.

Lemma 2. *Fix $\mathbf{p} = (p_0, p_1, \dots, p_n) \in \Delta(\{0, 1, \dots, n\})$. Any optimal search rule is such that:*

1. *Boxes are opened in decreasing order of p_i/c_i .*
2. *Whenever $B \subseteq [n]$ is a maximal class of boxes satisfying $p_i/c_i = \lambda > 0$ for all $i \in B$, the rule either does not enter B or searches it exhaustively.*

The resemblance of part 1 to Weitzman's policy may seem surprising given the strong negative dependence induced by Nature's one-prize distribution. The source—and limits—of this similarity become clearer once Weitzman's policy is decomposed into a selection and a stopping rule. The selection rule survives: conditional on continuing search, the DM opens boxes in decreasing order of the probability-to-cost ratio p_i/c_i , or equivalently in decreasing order of the binary Weitzman reservation value $z_i = \bar{u} - c_i/p_i$ when $p_i > 0$. To see why, compare two boxes i and j . Searching i before j is beneficial in the state where i contains the prize, because it saves the cost of opening j . Searching j before i is beneficial in the state where j contains the prize, because it saves the cost of opening i . Hence i should precede j exactly when

$$p_i c_j \geq p_j c_i,$$

or equivalently when $p_i/c_i \geq p_j/c_j$.

The stopping rule does not survive, however. Since rewards are correlated in the reduced problem, opening a box has experimentation value: finding no reward in one box makes it more likely that the prize is hidden elsewhere. As a result, the DM may be willing to open a box even when opening it and then stopping would yield a negative expected payoff. This feature is reflected in part 2 of [Lemma 2](#). If a group of boxes has the same probability-to-cost

ratio, the optimal rule treats the group as a unit: the DM either stops upon reaching the block or searches the block exhaustively. Searching only part of the block may yield a loss on its own, but this loss can be compensated by the information that an unsuccessful search provides about the remaining boxes in the block. The same principle applies across blocks. Exploring a block is optimal if the likelihood of finding the prize in that block, together with the continuation value generated by not finding it there, justifies the cost of exhaustively searching the block. Thus, the Weitzman index survives as a selection device, but not as a box-by-box stopping index.

The second step of the proof shows that, at a saddle point, these probability-to-cost ratios must be ordered by the cost ranking.

Lemma 3. *Let $(a, \mathbf{p})$ be a saddle point of the reduced game. Then,*

$$\frac{p_i}{c_i} \geq \frac{p_j}{c_j} \quad \text{for all } 1 \leq i < j \leq n.$$

Suppose instead that $p_i/c_i < p_j/c_j$ for some cheaper box i and more expensive box j . By [Lemma 2](#), the DM searches j before i . Conditional on this search rule, the DM would do weakly worse if the prize is in i than it is in j , because the rule reaches i later. Moreover, the oracle does strictly better in state i , since $c_i < c_j$. Regret is therefore strictly higher when the prize is in i than when it is in j . Nature could then profitably move probability from state j to state i , contradicting optimality.

The structure theorem follows by combining these properties. [Lemma 1](#) gives the reduced one-prize representation. [Lemma 3](#) implies that if a more expensive box is active, then every cheaper box must also be active; hence the active prize states form an initial segment. The maximal sets on which p_i/c_i is constant are contiguous blocks, and [Lemma 2](#) implies that the DM searches these blocks in order, treating each block as an all-or-nothing unit.

Saddle-point requirements. [Theorem 1](#) gives the architecture of any saddle point but leaves open the question of *how exactly* boxes are partitioned under the robustly optimal search rule. Starting from a block profile, the two support conditions of the zero-sum game must be satisfied: every state used by Nature must generate the same maximal regret, and every search path used by the DM must be optimal under Nature’s assignment of rewards. Operationally, the DM’s continuation probabilities between blocks equalize regret across active blocks, her lottery over within-block orders equalizes regret across prize locations within a block, and Nature’s reward distribution makes the DM willing to follow the prescribed block order and indifferent exactly at histories where she randomizes between stopping and continuing. While we postpone the formal treatment of these requirements to [Section 3.3](#),

let us explain the high-level idea here.

□ *Regret equalization.* To hedge against uncertainty over where the prize is hidden, the DM’s search rule equalizes regret *across* and *within* all active blocks. To gain intuition for which boxes are grouped and which are separated, consider two consecutive boxes i and j with $c_i < c_j$. The cheaper box i yields a higher oracle payoff than j , since $\bar{u} - c_i > \bar{u} - c_j$. Nature’s indifference requires that j ’s disadvantage in oracle payoff be compensated by a higher expected search cost when Nature hides the prize in j . Hence, in expectation, the cheaper box i must be searched before the more expensive box j . If the cost differential between the two boxes is sufficiently large, the required expected-search-cost difference cannot be implemented by a lottery over orders within one block. Box i must then be placed deterministically before j , so the two boxes belong to separate blocks. If, instead, costs are close to each other, the oracle payoffs are similar, and a deterministic ranking would create too large a difference in cumulative search costs across the two prize states. The DM must therefore randomize between opening box i or box j first, so the two boxes are placed in the same block.

□ *Worst-case belief.* Optimality of the block-search rule for the DM requires that the probability-to-cost ratios $(\lambda^{(k)})_{k=1}^K$ are strictly decreasing in k and that, upon reaching an active block $B^{(k)}$, the DM weakly prefers entering it over opting out—with exact indifference when her strategy prescribes randomized stopping upon reaching $B^{(k)}$. Hence, conditional on reaching any active block, the chance of finding the prize must be sufficiently high to justify the average search cost. This average search cost is formalized as follows. For each $S \subseteq [n]$, let

$$\Phi(S) := \frac{C(S)^2 + \Sigma(S)}{2C(S)}, \quad \text{where } C(S) := \sum_{i \in S} c_i \text{ and } \Sigma(S) := \sum_{i \in S} c_i^2.$$

The object $\Phi(S)$ has a simple interpretation. Suppose the prize is known to lie somewhere in S and Nature’s probability-to-cost ratio is constant on S , so that the conditional probability of box i is $c_i/C(S)$. Then $\Phi(S)$ is the expected cumulative cost of locating the prize when the DM searches S exhaustively. A key feature of this expectation is that it is independent of the exhaustive order used by the DM.⁹ Optimality of entering block $B^{(k)}$ then requires that the conditional expected reward outweigh the average cost of exhaustively searching $B^{(k)}$, given by $\Phi(B^{(k)})$. If $\Phi(B^{(k)})$ is sufficiently small, $\lambda^{(k)}$ can be calibrated to keep the DM indifferent between stopping and entering block $B^{(k)}$.¹⁰ Since $\Phi(B^{(k)})$ increases as boxes are

⁹The formal accounting identity is stated and proved in [Lemma A.1](#). In particular, for any exhaustive order of S , if E_i denotes the cumulative cost incurred when the prize is in box i , then $\sum_{i \in S} c_i E_i = C(S)\Phi(S)$. Dividing by $C(S)$ gives the stated cost-proportional expectation.

¹⁰Of course, the aggregate prize probability across active blocks cannot exceed one, which is why some

added to $B^{(k)}$, this limits how many boxes can be grouped together under the saddle-point block profile.

Global profitability. Applied to the whole set $[n]$, the term $\Phi([n])$ has a useful interpretation: conditional on the prize being somewhere, $\Phi([n])$ is the average cost of exhaustive search when all boxes have the same probability-to-cost ratio, so there is no inference that allows the DM to target more promising options first. Now suppose

$$\bar{u} \geq \Phi([n]). \quad (2)$$

Under this “global profitability” condition, exhaustive search remains worthwhile even in the least informative case, provided a prize is known to exist. We next show that this condition, and its violation, have direct implications for the properties of the saddle-point block profile. To state the result, we abstract from uninteresting boundary cases and impose the following genericity condition:

$$\bar{u} \neq \Phi(S) \quad \text{for every nonempty } S \subseteq [n]. \quad (3)$$

Proposition 1. *Suppose genericity condition (3) holds.*

1. *If $\bar{u} > \Phi([n])$, then, at every saddle point, Nature’s support is*

$$\{0, 1, \dots, n\},$$

and the DM inspects every box with positive probability.

2. *If $\bar{u} < \Phi([n])$, then every saddle point has at least two active blocks.*

Proposition 1 gives the sense in which $\Phi([n])$ disciplines the saddle-point structure. Suppose first that $\bar{u} > \Phi([n])$. Then, even if Nature gives the DM no guidance within the menu—in the sense that the probability-to-cost ratio is constant across all boxes—the full menu is still worth searching when the prize is known to be present. To keep the DM willing to randomize over whether to begin, Nature must therefore put positive probability on the no-prize state. Once the no-prize state is in support, every prize location must also be in support; otherwise Nature could move the prize to an ignored box and strictly raise regret. The worst-case belief thus has full support, and the DM’s optimal search rule inspects every box with positive probability.

boxes may have to be ignored by both the DM and Nature.

The case $\bar{u} < \Phi([n])$ is different. If all boxes were pooled into a single undifferentiated block, then exhaustive search of the whole menu would not be profitable in the cost-proportional benchmark. A single block would therefore provide too little direction: the DM would either want to stop immediately or, if induced to search on a subset of boxes, Nature would have a profitable deviation to boxes outside the block. Any saddle point must instead split the active boxes into at least two blocks. The point is not that adjacent costs must be far apart. Rather, when the menu as a whole is too costly relative to $\bar{u}$, Nature can sustain search only by giving the DM some direction about where to begin and by allowing her to stop between groups of boxes.

3.2 Tail-Block Search under Concave Costs

Before moving to the formal treatment of the saddle-point requirements laid out in the previous subsection, we consider a shape restriction on the cost profile that gives us a sharp characterization. In particular, in addition to the maintained ordering $c_1 < \dots < c_n$, we assume that cost differences are weakly decreasing along the sequence:

$$c_{i+1} - c_i \leq c_i - c_{i-1} \quad \text{for all } i = 2, \dots, n-1. \quad (4)$$

We will show that under this restriction the saddle point has a tail-block structure:

$$\{1\}, \{2\}, \dots, \{\ell-1\}, [\ell, n].$$

The economic intuition for this pattern follows from a simple saddle-point tradeoff. Consider two consecutive boxes, i and $i+1$. Equalizing regret requires the marginal oracle advantage of box i to be offset by the higher search priority assigned to it. The magnitude of the oracle effect is determined by the cost difference $c_{i+1} - c_i$, whereas the magnitude of the priority effect is determined by c_i , the cost wasted when the DM searches box i in state j . Concavity weakens the first force relative to the second as costs rise: the adjacent oracle advantage, $c_{i+1} - c_i$, weakly decreases, whereas the regret generated by searching the earlier box in vain, c_i , strictly increases. This suggests using a deterministic ranking among low-cost boxes, where the priority effect remains weak, and randomization in the upper tail, where a deterministic ordering would give Nature too strong an incentive to place the prize in a later box.

The theorem below establishes that this local force can be reconciled with the remaining saddle-point conditions. The singleton blocks are searched in increasing order of costs, with the DM potentially stopping between them. Once the DM reaches the tail, a fixed order

would create too much regret imbalance across possible prize locations. The DM therefore randomizes over orders within the tail and, conditional on entering it, searches the tail exhaustively.

To state the result, write

$$C_{\ell:n} := \sum_{i=\ell}^n c_i, \quad \Sigma_{\ell:n} := \sum_{i=\ell}^n c_i^2.$$

Theorem 2. *Assume concave costs and genericity condition (3). There exists a saddle point whose active prize states are $[n]$ and whose block partition is*

$$\{1\}, \{2\}, \dots, \{\ell^* - 1\}, [\ell^*, n],$$

where

$$\ell^* := \min \left\{ \ell \in [n] : \bar{u} > \Phi([\ell, n]) \text{ and } \bar{u} \geq \bar{u}_{\ell:n} \right\}, \quad (5)$$

and

$$\bar{u}_{\ell:n} := c_\ell + \frac{2C_{\ell:n}^2(\Sigma_{\ell:n}/C_{\ell:n} - c_\ell)}{C_{\ell:n}^2 - \Sigma_{\ell:n}}, \quad \ell < n; \quad \bar{u}_{n:n} := c_n. \quad (6)$$

The theorem turns the abstract block structure into an explicit search procedure. The DM begins with the low-cost boxes, which are searched one by one in increasing order of cost. After an unsuccessful search of a singleton block, she may stop before entering the next block. This is the region of *hesitant exploration*: the menu is still large, and the DM is concerned both about searching in vain and about searching the wrong part of the menu. Because cost differences provide relatively sharp guidance among low-cost boxes, these boxes provide a natural direction for search, but the DM still randomizes over how far to continue.

If the tail block is reached, the nature of the problem changes. The remaining boxes are relatively more similar, and the menu has become more manageable. Conditional on entering the tail, the DM searches it exhaustively, using a randomized order within the block. Thus the same robust rule that generates hesitation early in the search path also generates full exploration toward the end: once the DM has narrowed the problem to a sufficiently homogeneous tail, the value of continuing dominates the cost of protecting against further selection mistakes.

The threshold ℓ^* is the first tail boundary for which two requirements are met. The profitability requirement, $\bar{u} > \Phi([\ell, n])$, says that, once the DM reaches the tail and the prize is known to lie somewhere in it under Nature's one-prize distribution, the prize is valuable enough to justify exhaustive tail search. The second requirement, $\bar{u} \geq \bar{u}_{\ell:n}$, implies that regret equalization within the tail block is implementable by some stochastic search

order over the tail. Why implementability reduces to a single inequality on $\bar{u}$ is explained in the next subsection.

3.3 Saddle-Point Requirements: Formal Conditions

This subsection provides the verification machinery behind the block construction. This machinery is useful for checking candidate block profiles for arbitrary cost structures and for proving [Theorem 2](#), but it is not needed for the main economic message derived below. Readers interested primarily in that message can skip to [Section 4](#) on a first reading.

Fix a candidate partition of the active boxes into blocks $B^{(1)}, \dots, B^{(K)}$. By [Theorem 1](#), Nature's probabilities must satisfy $p_i/c_i = \lambda^{(k)}$ for all $i \in B^{(k)}$, with $\lambda^{(1)} > \dots > \lambda^{(K)}$ chosen to make the DM indifferent between continuing and stopping at the start of every non-terminal block. The DM's strategy can be described by two objects. The first is a sequence of reach probabilities $Q^{(1)}, \dots, Q^{(K)}$, where $Q^{(k)}$ is the probability that the DM enters block $B^{(k)}$ in a state in which no earlier block contains the prize. The second is, for each non-singleton block, a lottery over the deterministic orders of boxes within that block. Reach probabilities are calibrated to make Nature indifferent across blocks; the lottery over search orders is chosen to make Nature indifferent within a block.

DM's stopping indifference. Write the probability that the prize is in $B^{(k)}$ as

$$P^{(k)} = \sum_{i \in B^{(k)}} p_i = \lambda^{(k)} C(B^{(k)}).$$

The DM's expected payoff from exhaustively searching block $B^{(k)}$ and then stopping, conditional on the prize not being in one of the earlier blocks, is

$$V^{(k)} := \frac{P^{(k)}}{p_0 + \sum_{\ell=k}^K P^{(\ell)}} (\bar{u} - \Phi(B^{(k)})) - \frac{p_0 + \sum_{\ell=k+1}^K P^{(\ell)}}{p_0 + \sum_{\ell=k}^K P^{(\ell)}} C(B^{(k)}).$$

The first term captures the event where the prize lies in block k , which implies that the DM obtains $\bar{u}$ and pays the average cost $\Phi(B^{(k)})$. The second term captures the opposite case—there is no prize or it is located in a later block—so the DM pays the cost of searching the entire block without finding the prize.

We distinguish two cases according to whether the DM randomizes at the start of every block—this occurs when $p_0 > 0$ —or only at the start of every non-terminal block—this occurs when $p_0 = 0$. If the DM opts out with positive probability at the start of every active block, the continuation value is always zero, so $V^{(k)} = 0$ for all $k = 1, \dots, K$. These

indifference conditions together with $p_0 + \sum_k P^{(k)} = 1$ pin down the vector $(P^{(k)})_{k \in \{1, \dots, K\}}$ and p_0 . One then has to verify that this solution is a valid probability vector and that the resulting ratios

$$\lambda^{(k)} = \frac{P^{(k)}}{C(B^{(k)})}.$$

are indeed declining in k . A necessary condition for the solution to be a valid probability system, in particular for $P^{(k)} \geq 0$ for all k , is that each block is sufficiently inexpensive as a group, so that $\bar{u} > \Phi(B^{(k)})$ for all k ; that is, if the prize were guaranteed to be in the block, the reward would justify the average cost of finding it.

The other case is where, conditional on reaching it, the DM enters the final block $B^{(K)}$ with probability one and $p_0 = 0$. Conditional on having eliminated earlier blocks as possible prize locations, the DM is then certain to find the prize and obtains payoff

$$V^{(K)} = \bar{u} - \Phi(B^{(K)}) > 0.$$

To offset the positive continuation value, $V^{(K-1)}$ must now be negative.¹¹ For all preceding blocks $k < K - 1$, the payoff $V^{(k)}$ continues to equal zero.

Nature's across-block indifference. Just as Nature's prize probabilities keep the DM indifferent between exiting and continuing between blocks, the DM's continuation probabilities between blocks must be calibrated to make Nature indifferent between assigning the prize to any given block and, potentially, not assigning a prize at all. Given the DM's reach probabilities, the average regret generated by hiding the prize in block $B^{(k)}$ is

$$R^{(k)} = (1 - Q^{(k)}) \left(\bar{u} - \frac{\Sigma(B^{(k)})}{C(B^{(k)})} \right) + \sum_{m < k} Q^{(m)} C(B^{(m)}) + Q^{(k)} \left(\Phi(B^{(k)}) - \frac{\Sigma(B^{(k)})}{C(B^{(k)})} \right). \quad (7)$$

The first term is the oracle payoff forgone when the DM stops before reaching block $B^{(k)}$ —the reward $\bar{u}$ minus the expected cost of finding the prize immediately, which, given Nature's random prize allocation, equals

$$\sum_{i \in B^{(k)}} \frac{c_i}{C(B^{(k)})} \cdot c_i = \frac{\Sigma(B^{(k)})}{C(B^{(k)})}.$$

The second term is the cost wasted on earlier blocks, which do not contain the prize in this state. The third term is the within-block selection cost: conditional on entering the block,

¹¹This does not imply that the DM would skip $B^{(K-1)}$: because $\lambda^{(K-1)} > \lambda^{(K)}$, [Lemma 2](#) requires the former block to be searched first.

the DM finds the prize but pays the average cumulative cost $\Phi(B^{(k)})$, while the oracle pays only the average cost of directly opening the “right” box, $\Sigma(B^{(k)})/C(B^{(k)})$. This last term is zero for singleton blocks.

At a saddle point, $R^{(k)}$ must be the same for all active blocks. The no-prize state generates regret

$$R^{(0)} = \sum_{m=1}^K Q^{(m)} C(B^{(m)}). \quad (8)$$

This value must be no larger than the common active-prize regret, and it is equal to that regret whenever Nature assigns positive probability to the no-prize state. These equalities and inequalities determine the reach probabilities.

Nature’s within-block indifference. Finally, the DM’s selection rule within blocks must equalize Nature’s regret across all boxes contained in a given block. To formalize this requirement, fix an active block $B^{(k)}$ and its reach probability $Q^{(k)} > 0$. Let E_i be the expected cumulative cost incurred inside the block under the DM’s lottery over exhaustive search orders, conditional on entering the block and conditional on the prize being in box $i \in B^{(k)}$. Equalizing regret between two boxes $i, j \in B^{(k)}$ requires:

$$c_j - c_i = Q^{(k)}(E_j - E_i). \quad (9)$$

The left-hand side represents the oracle’s absolute payoff advantage from finding the prize in the cheaper box i rather than the more expensive box j . For Nature to remain indifferent, this oracle advantage must be offset by the right-hand side, which captures the DM’s additional expected search costs from prioritizing the cheaper box i over j in the randomized search order when the prize is in fact in j . Condition (9), together with an accounting identity from [Lemma A.1](#) in the Appendix, pins down the regret-equalizing target vector¹²

$$E_i = \frac{c_i}{Q^{(k)}} + \Phi(B^{(k)}) - \frac{\Sigma(B^{(k)})}{Q^{(k)}C(B^{(k)})}, \quad i \in B^{(k)}. \quad (10)$$

The implementability problem. Equation (10) gives the expected cumulative search costs that would render Nature indifferent across all boxes in a block. The key question is feasibility: can some lottery over within-block search orders generate this vector $(E_i)_{i \in B^{(k)}}$? The following result gives the answer.

Proposition 2. *A vector $E = (E_i)_{i \in B}$ is implementable by a lottery over permutations of B*

¹²See [Appendix A.5](#) for details.

if and only if

$$\sum_{i \in S} c_i E_i \geq \frac{C(S)^2 + \Sigma(S)}{2} \quad \text{for every } S \subseteq B, \quad (11)$$

with equality at $S = B$.

The inequality has a simple interpretation. The term

$$v(S) := \frac{C(S)^2 + \Sigma(S)}{2}$$

is the smallest cost-weighted cumulative search cost that the boxes in S can generate: it is attained when all boxes in S are searched before any box outside S . Therefore, no lottery over orders can assign the boxes in S a lower total cost-weighted cumulative search cost than $v(S)$. This proves necessity. When $S = B$, the inequality is an equality because the total cost-weighted cumulative search cost of exhaustively searching the whole block is invariant to the order.

For sufficiency, consider a hypothetical *cooperative game* with *characteristic function* $v(S)$. The inequalities in (11), together with equality at B , say that the vector $(c_i E_i)_{i \in B}$ lies in the core of this game. Moreover, the marginal contribution of any box j ,

$$v(S \cup \{j\}) - v(S) = c_j (C(S) + c_j),$$

becomes larger when the coalition S that is already present becomes larger. The game is therefore convex, or supermodular. By the core characterization for convex games—equivalently, the polymatroid characterization—the core is the convex hull of its extreme points, each consisting of marginal contributions of boxes along a sequence of nested subsets within B (Shapley, 1971; Edmonds, 2003; Ichiishi, 1981). Each sequence of nested subsets induces a permutation of boxes in B , and, importantly, the associated marginal contributions coincide with the DM's cost-weighted cumulative search costs when she searches according to that order.¹³ This means that if the regret-equalizing cost-weighted expenditure vector $(c_i E_i)_{i \in B}$ lies in the core, satisfying (11), it can be implemented by a lottery over deterministic search orders. We thus conclude that (11) guarantees the implementability of E via some randomized search order.

¹³To see this, let $\pi = (\pi_1, \dots, \pi_m)$ be the induced permutation and let $S_k = \{\pi_1, \dots, \pi_k\}$. Searching according to π yields cumulative cost $E_{\pi_k}^\pi = C(S_k)$ when π_k is the prize box, while the corresponding marginal contribution is

$$v(S_k) - v(S_{k-1}) = c_{\pi_k} C(S_k) = c_{\pi_k} E_{\pi_k}^\pi.$$

Hence the marginal vector associated with the nested sequence is exactly the cost-weighted cumulative-search-cost vector $(c_i E_i^\pi)_{i \in B}$ generated by the order π .

The core test in [Proposition 2](#) requires checking every subset of B . The target vector in (10), however, has additional structure. Because the weighted target $c_i E_i$ contains a positive quadratic term in c_i and a common linear term, coalitions that contain expensive boxes receive relatively high weighted cumulative costs. The constraints that can bind are therefore the lower-cost coalitions. To formalize this, write the boxes in B in increasing cost order as $b_1, \dots, b_{|B|}$, and define the prefix sets

$$S_k := \{b_1, \dots, b_k\}, \quad k = 1, \dots, |B|.$$

Proposition 3. *Fix a block B and the target vector $(E_i)_{i \in B}$ defined in (10). Then (11) holds for all sets $S \subseteq B$ if and only if it holds for all prefix sets $S_k \subseteq B$.*

[Proposition 3](#) shows that implementability is hardest for the cheap prefixes. These boxes are precisely the ones that the target tries to reach early. If the lottery over search orders can give enough priority to every cheap prefix, then coalitions containing more expensive boxes automatically satisfy the core inequalities. The high-dimensional implementability problem thus collapses to a nested sequence of prefix constraints.

In the case of concave costs, the simplification goes one step further. Not only do we need to check only prefixes; among the prefixes of a tail, the binding one is the singleton prefix consisting of the cheapest tail box.

Lemma 4. *Assume costs are concave and fix a tail block $T = [\ell, n]$. If T is a singleton, within-block implementability is automatic. If $\ell < n$, then within-block implementability is satisfied if and only if the core condition (11) holds for the singleton prefix $\{\ell\}$. Letting Q^T be the DM's reach probability of tail T , this condition can be written as*

$$Q^T \geq \frac{\Sigma(T)/C(T) - c_\ell}{\Phi(T) - c_\ell}. \quad (12)$$

In the case of concave costs, among all prefix constraints for T , the most demanding one is the singleton prefix $\{\ell\}$, the cheapest box in the tail. If the lottery over search orders can give sufficient priority to this cheapest box, i.e., if (12) holds, then concavity guarantees that all longer prefixes are also feasible. Noting that the DM's probability of entering the tail Q^T increases with the size of the potential reward $\bar{u}$, the implementability requirement reduces to the simple cutoff condition $\bar{u} \geq \bar{u}_{\ell:n}$.

4 Comparative Statics

We now turn to some of the questions motivating our analysis. Digital platforms have changed consumer search in two ways. On the one hand, they have made the inspection of any given option substantially cheaper. On the other hand, they have dramatically expanded the number of available options. It has been noted that this may lead to paradoxical outcomes, where, even though each individual search is easier, consumers may feel overwhelmed by the size of the menu and disengage from search altogether.

Our model provides a simple mechanism for this phenomenon. When the menu grows, or when inspection costs become similar, so that available options become harder to rank *ex ante*, the DM's concern over mistakes from exploring the wrong options becomes more severe, reducing her inclination to search at all. The comparative statics below make this force precise: focusing on the saddle-point construction under concave costs from [Theorem 2](#), we consider the DM's probability of initiating search, $Q^{(1)}$, under the robustly optimal search rule and study how it changes when more options become available and when options become more similar *ex ante*.

4.1 Adding options

We first ask what happens when the menu expands. In a standard Bayesian search model, adding an option to an existing menu cannot hurt the DM: she can always ignore it. In the robust problem, however, an additional option gives Nature a new possible location for the prize. Since the worst case is a one-prize environment, the added option increases the scope for selection mistakes. The DM must now hedge against a larger set of possible ways in which she may search the wrong boxes. The next result shows that this force unambiguously reduces the probability of initiating search.

Theorem 3. *Let $c = (c_1, \dots, c_n)$ be strictly increasing and concave, and suppose a new box with cost $c_{\text{new}} > c_1$ is inserted at position $k \in \{2, \dots, n+1\}$, yielding a strictly increasing concave sequence $\tilde{c} = (\tilde{c}_1, \dots, \tilde{c}_{n+1})$. Let $Q^{(1)}$ and $\tilde{Q}^{(1)}$ be the probabilities of initiating search in the saddle-point construction of [Theorem 2](#) under c and $\tilde{c}$, respectively. Then*

$$\tilde{Q}^{(1)} \leq Q^{(1)}.$$

[Theorem 3](#) captures a stark form of choice overload. As the choice set grows, the DM becomes more likely to opt out before inspecting any option. This effect is driven not by the cost of inspecting each additional option in isolation, but by the increasing difficulty of navigating a larger menu: if there is only one prize, a larger menu provides more places for

it to be hidden and, consequently, more ways for the DM to begin searching in the wrong place.

The theorem concerns the extensive margin of search. Another natural variable of interest is the expected number of options inspected. In contrast to the initial search probability, this object depends not only on the robust search rule, but also on the true data-generating process (DGP) for rewards. [Figure 2](#) fixes this process to be i.i.d. Bernoulli, with probability p of a high reward, and plots the expected number of opened boxes as the menu size varies. Of course, not knowing this true DGP, the DM employs the robust search rule characterized earlier.




Figure 2: Expected number of opened boxes when rewards are i.i.d. Bernoulli with $p \in \{0, 0.1, 0.25, 0.5, 1\}$, as a function of n . For each n , costs vary between 0.1 and 0.2 with constant increments. The drops correspond to points at which the saddle-point structure transitions from a single block to multiple blocks.

The figure shows a hump-shaped pattern: for small menus, adding options can increase the total amount of search, while for large menus it decreases it. This is consistent with a standard theme in the choice-overload literature: additional options may initially be attractive, but sufficiently large assortments can reduce choice, purchase, and satisfaction ([Iyengar and Lepper, 2000](#); [Reutskaja and Hogarth, 2009](#)). Recent evidence from online recommender systems documents a closely related pattern: increasing the number of recommended products first raises and then lowers purchase, with much of the decline coming from a lower probability of starting search ([Long et al., 2025](#)). At the same time, the empirical literature emphasizes that choice overload is not a universal phenomenon; rather, its emergence

depends on factors such as choice-set complexity, decision difficulty, preference uncertainty, and the consumer’s decision goal (Scheibehenne et al., 2010; Chernev et al., 2015).

In our robust model, the hump shape comes from two opposing forces. As options are added, the DM becomes less likely to initiate search but, conditional on starting and finding nothing in the early boxes, may search further. For small menus, the continuation effect dominates; for large menus, the decline in initiation more than offsets it. Thus choice overload is not modeled as a psychological anomaly. It is a rational response to a harder robust-search problem.

4.2 Flattening the cost profile

The previous result shows that expanding the menu can reduce search. Our interpretation is that the key friction is the absence of guidance. To substantiate this reading of the model, we next consider the effect of changes in the cost profile. When costs differ sharply, the DM has a natural way to organize search: start with the cheap boxes, and only later move to the expensive ones. When the cost profile becomes flatter, this guidance weakens. The boxes become more similar from the point of view of search, and the DM has fewer reasons to choose one path over another.

The following comparative static formalizes this idea. It compares two concave cost profiles, where the second profile is both higher and flatter.

Theorem 4. *Let $c = (c_1, \dots, c_n)$ and $\tilde{c} = (\tilde{c}_1, \dots, \tilde{c}_n)$ be strictly increasing concave cost sequences such that $\tilde{c}_i \geq c_i$ for all i , and $\tilde{c}_i - \tilde{c}_{i-1} \leq c_i - c_{i-1}$ for all $i = 2, \dots, n$. Let $Q^{(1)}$ and $\tilde{Q}^{(1)}$ be the probabilities of initiating search in the saddle-point construction of [Theorem 2](#) under c and $\tilde{c}$, respectively. The probability of initiating search weakly decreases:*

$$\tilde{Q}^{(1)} \leq Q^{(1)}.$$

There are two forces in the theorem. The first is direct: higher search costs make search less attractive. The second is informational: a flatter profile makes it harder to use costs as a guide. It is important to emphasize that higher search costs alone do not guarantee a lower probability of initiating search. Indeed, it is easy to construct pairs of cost profiles in which one has uniformly lower but flatter costs, yet induces a lower probability of initiating search. A stark illustration arises in the homogeneous limit below, where costs become entirely flat. Whereas a steep profile gives the DM direction, a flat profile deprives her of guidance, leaving her with a more symmetric and therefore more ambiguous navigation problem.

Homogeneous limit. The cleanest way to isolate the role of flattening is to consider the limiting case in which cost differences vanish. Fix $c \in (0, \bar{u})$. For each flatness parameter $\nu > 0$, let

$$c^\nu = (c_1^\nu, \dots, c_n^\nu)$$

be a strictly increasing, concave cost profile satisfying $\max_{i \in [n]} |c_i^\nu - c| \leq \nu$. For each ν , let $(a_\nu, \mathbf{p}_\nu)$ be the saddle point constructed in [Theorem 2](#), and let ℓ_ν be its tail boundary. Thus the blocks are

$$B_\nu^{(1)} = \{1\}, \dots, B_\nu^{(\ell_\nu-1)} = \{\ell_\nu - 1\}, \quad B_\nu^{(\ell_\nu)} = [\ell_\nu, n],$$

with the convention that there is only the tail block when $\ell_\nu = 1$. Let $Q_\nu^{(k)}$ denote the probability of reaching block $B_\nu^{(k)}$. We study the limit as $\nu \downarrow 0$.

For a flat profile, the global profitability condition simplifies substantially:

$$\bar{u} > \Phi([n]) = \frac{n^2 c^2 + n c^2}{2 n c} = \frac{n+1}{2} c.$$

Fixing $\bar{u}$ and c , we observe that this condition becomes violated when n becomes sufficiently large. Let

$$\bar{n} := \max \left\{ k \in [n] : \frac{k+1}{2} c \leq \bar{u} \right\}$$

be the largest size of a homogeneous menu that can be searched profitably when the prize is known to be present.

Theorem 5. Fix $c \in (0, \bar{u})$. For each $\nu > 0$, let $c^\nu = (c_1^\nu, \dots, c_n^\nu)$ be a strictly increasing, concave cost profile satisfying

$$\max_{i \in [n]} |c_i^\nu - c| \leq \nu.$$

Let ℓ_ν be the tail boundary in the saddle-point construction of [Theorem 2](#), and let $Q_\nu^{(k)}$ be the probability of reaching the k th block in that construction. Then, as $\nu \downarrow 0$:

(i) If $\bar{u} > \frac{n+1}{2} c$, then $\ell_\nu = 1$ for all sufficiently small ν , and

$$Q_\nu^{(1)} \rightarrow Q^* := \frac{\bar{u} - c}{\bar{u} + \frac{n-1}{2} c} \in (0, 1).$$

(ii) If $\bar{u} < \frac{n+1}{2} c$, then

$$Q_\nu^{(k)} \rightarrow 0 \quad \text{for every } k = 1, \dots, n - \bar{n} + 1.$$

In particular, $Q_\nu^{(1)} \rightarrow 0$.

The theorem has two regimes. If $\bar{u} > \frac{n+1}{2}c$, the whole flat menu is globally profitable. In the limit, Nature assigns equal mass to all prize states and positive mass to the no-prize state. The DM initiates search with probability $Q^* = \frac{\bar{u}-c}{\bar{u}+\frac{n-1}{2}c}$ and, conditional on searching, explores all options. In the approximating constructions, $\ell_\nu = 1$ eventually, so the whole menu is treated as a single block. In the limit, Nature and the DM treat boxes symmetrically: each box has the same probability of containing the prize, and, conditional on searching, the DM uses a uniformly random order. Consistent with [Theorem 3](#), the probability of initiating search strictly declines with n : a larger homogeneous menu makes search less attractive because it increases the expected cost of finding the single prize.

If $\bar{u} < \frac{n+1}{2}c$, global profitability fails. In this regime, the probability of initiating search converges to zero as boxes become homogeneous. This is the starkest form of choice overload in the model: as the cost profile becomes flat, the DM disengages from search altogether. The implication is that lower search costs need not induce more search. Fix any strictly increasing concave profile $\hat{c}$ satisfying the maintained assumptions and $\hat{c}_i > c$ for every i , and let $\hat{Q}^{(1)} > 0$ be the initiation probability generated by [Theorem 2](#) under $\hat{c}$. Then, for all sufficiently small ν , the profile c^ν is uniformly cheaper than $\hat{c}$, yet $Q_\nu^{(1)} < \hat{Q}^{(1)}$. Thus the DM searches more under the uniformly higher, but more informative, cost profile than under the uniformly lower, sufficiently flat one.

The limiting belief in the second regime is revealing: Nature places no mass on the no-prize state ($p_0^* = 0$). The DM's disengagement is therefore not driven by the possibility that there is nothing to find, but by the difficulty of locating the prize in a large, homogeneous menu. Under the worst-case scenario, the prize is always hidden somewhere, yet the DM has no indication of where to search first. When the menu is sufficiently large relative to the value of the prize, this absence of direction is enough to eliminate search.

The homogeneous limit therefore sharpens the interpretation of the previous comparative statics. Adding options reduces search because it expands the set of possible mistakes. Flattening the cost profile reduces search because it removes the cost-based ranking that helps the DM avoid those mistakes. Digital platforms combine both forces: they make individual inspection cheap, but they also create large and often weakly differentiated menus. The robust-search friction is precisely that, in such environments, the consumer's problem is not only whether to search, but how to begin.

5 Discussion and Robustness

We conclude by examining three robustness issues. The binary-reward assumption is the clearest limitation of the baseline model: although it yields a clean characterization, it leaves

open the extent to which the overload result is driven by this simplification. We address this concern first. We then return to the binary benchmark to examine two further dimensions: dynamic reoptimization and restricting Nature’s ambiguity set to i.i.d. reward distributions.

To streamline the exposition, we restrict attention to the homogeneous case in which $(c_1, \dots, c_n) = (c, \dots, c)$, with $c \in (0, \bar{u})$. This benchmark corresponds to the flat limit characterized in [Theorem 5](#). In particular, if global profitability fails—that is, if $\bar{u} < \frac{n+1}{2}c$ —the DM opts out with probability one. By contrast, if $\bar{u} > \frac{n+1}{2}c$, the robustly optimal commitment rule initiates search with probability

$$Q^*(n) = \frac{\bar{u} - c}{\bar{u} + \frac{n-1}{2}c}, \quad (13)$$

where we make the dependence on n explicit. Conditional on initiating search, the DM searches exhaustively, selecting each remaining box uniformly at random at every stage.

5.1 Beyond Binary Rewards

The binary-reward assumption is the main tractability device in the baseline analysis: it renders the state space finite and makes stopping automatic once a high reward is found. We now relax this restriction and set

$$\Omega = [0, \bar{u}]^n,$$

with Nature choosing any distribution from $\Delta([0, \bar{u}]^n)$.¹⁴

The DM’s strategy must now condition not only on the set of remaining boxes, but also on the rewards observed so far. In particular, the current best reward $z \in [0, \bar{u}]$ becomes a state variable: after stopping, the DM takes the best currently available payoff, including the outside option. The following simple observation allows us to focus on rules that treat the unopened boxes symmetrically.

Lemma 5. *Suppose $\Omega = [0, \bar{u}]^n$ and costs are homogeneous. If the robust search problem has a saddle point, then it has one in which Nature’s distribution is exchangeable and the DM’s rule is symmetric: at every history, conditional on continuing search, each unopened box is selected with equal probability.*

The lemma removes labels from the choice of which unopened box to inspect. The description of the search rule can thus be reduced to a continuation probability as a function of the previously observed rewards.

¹⁴Without the upper bound $\bar{u}$, the minmax regret problem need not have a finite solution.

For a sufficiently small upper bound $\bar{u}$, the DM employs a threshold rule. In particular, we show in [Appendix B.7.2](#) that, in the two-box problem, after opening the first box and observing reward z , she opens the second box if $z < \bar{u} - c$ and stops otherwise. When the upper bound $\bar{u}$ is large, however, a fixed acceptance threshold exposes the DM to two opposite forms of regret. A low threshold risks stopping too early after a mediocre draw, while a high threshold risks excessive search after a sufficiently good draw. The optimal rule then requires a randomized acceptance threshold (see [Appendix B.7.2](#)).

The additional state dependence makes the general problem substantially less tractable than the binary model. The main behavioral implication, however, is unchanged: sufficiently large homogeneous menus lead to complete disengagement.

Corollary 1. *Let $\Omega = [0, \bar{u}]^n$, and suppose search costs are homogeneous. If*

$$\bar{u} \leq \frac{n+1}{2}c,$$

then immediate opt-out is robustly optimal.

The proof uses the binary model as a lower bound on Nature’s power. Restrict Nature to the binary one-prize distributions used in the homogeneous benchmark. When $\bar{u} \leq (n+1)c/2$, the binary problem already allows Nature to impose regret $\bar{u} - c$, which is exactly the regret from immediate opt-out. Enlarging Nature’s feasible set from binary distributions to all distributions on $[0, \bar{u}]^n$ cannot reduce the maximal regret of any search rule. At the same time, opting out immediately guarantees regret of at most $\bar{u} - c$, since the oracle’s net payoff cannot exceed $\bar{u} - c$. Hence immediate opt-out remains minmax optimal in the richer problem.

The result shows that choice overload is not an artifact of binary rewards. Even when rewards can take a continuum of values, a sufficiently large homogeneous menu makes search unattractive.

5.2 Dynamic Reoptimization

It is well known that a choice based on the minmax regret criterion can lead to *dynamic inconsistency*, arising from two potential sources (see [Hayashi \(2011\)](#) for details): (1) new information may change the regret-maximizing scenario guiding the DM’s choice; (2) regret-based preferences are opportunity dependent—removing some alternatives can change the optimal choice among the remaining ones.

In the homogeneous benchmark, this issue takes a simple form. When global profitability holds so that search is worthwhile, the commitment solution randomizes once at the outset;

conditional on opening the first box, the DM plans to search exhaustively until she finds the prize or exhausts the menu. If, after an empty draw, she is allowed to reoptimize, she instead views the remaining problem as a fresh $(n - 1)$ -box problem. The saddle-point belief for this continuation problem is not the Bayesian update of the original n -box saddle-point belief: it assigns less probability to the remaining prize states, or equivalently more probability to the no-prize state. Thus, as the DM eliminates options without finding a prize, she becomes increasingly concerned that no prize exists and has an incentive to re-randomize rather than carry out the original exhaustive-search plan.

Let us consider a *sophisticated* DM who anticipates this future behavior but cannot commit to a plan. Her problem can be viewed as an *intrapersonal game*, in which the DM's current self chooses an action at the current stage, taking future selves' continuation behavior as given. For each set $N \subseteq [n]$ of remaining boxes, let

$$\Omega_N := \{\omega \in \Omega : u_i = 0 \text{ for all } i \notin N\}$$

denote the set of states consistent with having reached N without discovering a high reward. Let $u_N(a, \omega)$ denote the continuation payoff, with past search costs sunk, from following strategy a in the subproblem with remaining set N .

Definition 1. A strategy $a = (a_N)_{N \in \mathcal{N}}$ is an **intrapersonal equilibrium** if, for each set $N \in \mathcal{N}$ of remaining boxes,

$$a_N \in \arg \min_{\tilde{a}_N \in \Delta(N \cup \{0\})} \max_{\omega \in \Omega_N} \left[u_N(a^*(\omega), \omega) - u_N((\tilde{a}_N, (a_{N'})_{N' \subsetneq N}), \omega) \right].$$

An intrapersonal equilibrium can equivalently be interpreted as a *sequential saddle point*: at every continuation problem, the current self chooses a stage action that minimizes continuation regret against Nature's worst-case state, taking as given the strategy followed by future selves.

Theorem 6. *Assume the DM cannot commit and search costs are homogeneous. There exists an intrapersonal equilibrium and an integer cutoff*

$$\bar{n}^{nc} \leq \bar{n}$$

such that, at any history with m unopened boxes:

1. *if $m \leq \bar{n}^{nc}$, the DM randomizes between opening a uniformly selected remaining box and opting out;*

2. if $m > \bar{n}^{nc}$, the DM opts out with probability one.

The cutoff structure in [Theorem 6](#) reflects the tension between the current self’s incentive to search and her anticipation that future selves will reoptimize. To see this, consider first a naive DM who fails to foresee future incentives to reoptimize. Such a DM would adopt the n -box commitment plan at the initial node, initiating search with probability $Q^*(n)$, but would then renege after each empty draw and solve the smaller continuation problem anew, using probabilities $Q^*(n-1)$, $Q^*(n-2)$ and so on, rather than following the original exhaustive-search plan. A sophisticated DM anticipates this future tendency to under-search. To preserve her future chances of finding the prize, she may therefore choose a higher initial search probability than under the corresponding commitment plan. This force can make the initial search probability non-monotone in the number of options, in contrast to the commitment comparative static in [Theorem 3](#). The large-menu conclusion is nevertheless unchanged: once the menu is sufficiently large, even the sophisticated DM opts out with probability one.

The cutoff $\bar{n}^{nc}$ is weakly below the commitment cutoff $\bar{n}$ because the inability to commit makes initiating search less valuable. Indeed, choice overload has two sources now. The first is the baseline needle-in-a-haystack problem: adding options expands the set of possible locations in which the prize may be hidden. The second is intrapersonal: without commitment, the DM cannot ensure that future selves will carry out the search plan that made the initial search worthwhile.

We provide an explicit characterization of the equilibrium search probabilities and regret levels in [Appendix B.8](#).

5.3 Independent and Identically Distributed Rewards

The baseline model allows Nature to choose any distribution over reward assignments, including arbitrary correlation across boxes. We now consider a more structured ambiguity set. Rewards are independent and identically distributed: each box contains the high reward with probability $p \in [0, 1]$, independently across boxes, but the DM does not know the value of p . This is a robust analogue of the classical [Weitzman \(1979\)](#) environment.

By symmetry, conditional on searching, the DM can be taken to select an unopened box uniformly at random. Her dynamic strategy is therefore described by a vector

$$(\alpha_m)_{m \in \{1, \dots, n\}},$$

where α_m is the probability of continuing search when m boxes remain and no high reward

has yet been found. Given p , the minimized current-stage regret can be written as

$$R_m(p) := \min_{\alpha_m \in [0,1]} \left[(1 - \alpha_m)(1 - (1 - p)^m)(\bar{u} - c) + \alpha_m(1 - p)(c + R_{m-1}(p)) \right], \quad (14)$$

with $R_0(p) = 0$. The first term is the regret from stopping: if at least one of the remaining boxes contains the high reward, the oracle would have opened such a box and obtained $\bar{u} - c$. The second term is the regret from searching and drawing a low-reward box: the DM pays the cost c and then faces the $(m - 1)$ -box continuation problem.

The robustly optimal continuation rule is characterized as follows.

Theorem 7. *Assume rewards are independent and identically distributed and search costs are homogeneous. Conditional on not having found a high reward and on there being $m \in \{1, \dots, n\}$ unopened boxes, the DM opens a uniformly selected remaining box with probability*

$$\alpha_m^* = \frac{m(\bar{u} - c)^m}{(m - 1)(\bar{u} - c)^m + \bar{u}^m}.$$

Nature's worst-case success probability is

$$\hat{p} = \frac{c}{\bar{u}},$$

and the associated regret is

$$R_m^* = \left(1 - \left(\frac{\bar{u} - c}{\bar{u}} \right)^m \right) (\bar{u} - c).$$

The described strategy also constitutes the unique intrapersonal equilibrium.

The DM's regret-minimizing search behavior is illustrated in [Figure 3](#). For each fixed p , the expression in (14) is linear in α_m . For each fixed α_m , Nature chooses the success probability p that maximizes regret. The DM's worst-case regret is therefore the upper envelope of these linear functions, and the optimal search probability minimizes this envelope. In contrast to the baseline model with arbitrary correlation, the implementation of the optimal rule does not require commitment: the rule described in [Theorem 7](#) not only minimizes ex-ante regret but also constitutes an intrapersonal equilibrium.

The next result shows that, even when Nature is restricted to assigning rewards independently and identically—thereby effectively ruling out the possibility of hiding a single prize—adding options discourages the DM from initiating search.

Corollary 2. *Consider the solution described in [Theorem 7](#).*




Figure 3: Saddle point with $n = 3$ unopened boxes; $\bar{u} = 1$, $c = 0.3$.

1. The probability of search α_n^* is strictly decreasing in n , with

$$\lim_{n \rightarrow \infty} \alpha_n^* = 0.$$

2. The associated regret R_n^* is strictly increasing in n , with

$$\lim_{n \rightarrow \infty} R_n^* = \bar{u} - c.$$

As in the setting with arbitrarily correlated rewards, the model predicts choice overload. This demonstrates that, although the single-prize construction in the baseline model provides a transparent benchmark, it is not essential to the result: also with i.i.d. rewards, larger menus discourage the DM from initiating search. While a larger menu raises the probability that some option is valuable, it also increases the number of unsuccessful draws the DM may face before finding one. Under minmax regret, the latter force dominates in the limit: search probabilities converge to zero, and regret converges to the loss from missing a high-reward box altogether.

6 Conclusion

This paper revisits the classic Pandora's box problem by studying a DM who confronts ambiguity about the value of her options and aims to minimize her maximum anticipated regret. By doing so, we provide a tractable framework for analyzing sequential search in complex environments and offer a new theoretical foundation for the well-documented phenomenon

of choice overload.

Our analysis demonstrates that a regret-minimizing agent’s willingness to search is fundamentally shaped by the size of the choice set and the cost profile. The DM’s worst-case belief centers on a “hidden treasure” scenario, which generates a stark prediction: the DM’s propensity to search falls when new options are added to the choice set.

Crucially, our framework allows us to identify the precise mechanism driving this disengagement: selection error, or the fear of investing search effort in the wrong options. As the choice environment becomes more symmetric, the DM facing sufficiently many options becomes paralyzed by the ambiguity of where to start. This diagnosis points directly to potential remedies in *assortment design*: platforms may mitigate overload by deliberately breaking the symmetry of the menu—for instance, by offering recommendations or introducing cost heterogeneity. By *steering* the DM toward an “easy” path, such designs alleviate the regret associated with selection error and restore the incentive to explore.

A Proofs

A.1 Proof of **Lemma 1**

Fix $a \in \Delta(A)$. Let ω be a state with at least one prize, and let

$$i^*(\omega) \in \arg \min \{c_i : \omega_i = \bar{u}\}$$

be the cheapest prize location. Let ω' be the state that keeps only this prize:

$$\omega'_{i^*(\omega)} = \bar{u}, \quad \omega'_j = 0 \quad \text{for all } j \neq i^*(\omega).$$

The oracle payoff is the same in the two states, because the cheapest prize location is the same:

$$u(a^*(\omega'), \omega') = u(a^*(\omega), \omega).$$

It remains to compare the DM’s payoff. First, consider a pure plan $b \in A$. Since the set of prize boxes in ω' is a subset of the set of prize boxes in ω , the plan b cannot find a prize earlier in ω' than in ω . More precisely, up to the first time a prize is found in ω , the histories in ω and ω' coincide unless b opens a prize box that was removed in the construction of ω' . In that case, b stops in state ω , while in state ω' it has not found a prize and may only continue searching or stop later. Hence, relative to ω , the reward obtained by b in ω' weakly falls and the total search cost weakly rises. Therefore, $u(b, \omega') \leq u(b, \omega)$. Taking expectations over

the pure plans according to the mixed rule a gives

$$u(a, \omega') \leq u(a, \omega).$$

Together with equality of the oracle payoffs, this implies

$$u(a^*(\omega'), \omega') - u(a, \omega') \geq u(a^*(\omega), \omega) - u(a, \omega).$$

Now let $\pi \in \Delta(\Omega)$ be any regret-maximizing distribution. Construct a new distribution by leaving unchanged the probability assigned to states with no prize or exactly one prize, and by moving the probability assigned to each state with several prizes to the one-prize state constructed above. The preceding inequality shows that this operation weakly increases expected regret. Since π was already regret-maximizing, the new distribution is also regret-maximizing. By construction, it assigns positive probability only to states with at most one prize. $\square$

A.2 Proof of Lemma 2

For fixed $\mathbf{p}$, the oracle's expected payoff is independent of the DM's rule. Hence the DM's best replies to $\mathbf{p}$ are exactly the rules that maximize the DM's expected payoff under $\mathbf{p}$. Since this payoff is linear in the DM's mixed rule, it suffices to consider pure rules.

Under the reduced state space, a pure rule is identified with the finite list of boxes it opens along the history in which no prize is found. If a prize is found, the rule stops. After a set $S \subset [n]$ has been opened and found empty, the posterior is

$$\Pr(i \mid S) = \frac{p_i}{p_0 + \sum_{\ell \notin S} p_\ell} \quad \text{for every } i \notin S. \quad (15)$$

Step 1: Fix such a history S and compare two continuation plans that differ only by swapping the next two unopened boxes i and j . The two plans coincide after both boxes have been opened and found empty. Thus the conditional expected payoff difference from opening i before j , rather than j before i , is

$$\Pr(i \mid S)c_j - \Pr(j \mid S)c_i = \frac{p_i c_j - p_j c_i}{p_0 + \sum_{\ell \notin S} p_\ell}. \quad (16)$$

Hence opening i before j is weakly better if and only if

$$\frac{p_i}{c_i} \geq \frac{p_j}{c_j}.$$

Repeated adjacent swaps therefore produce an optimal rule that opens boxes in weakly decreasing order of p_i/c_i before stopping.

Step 2: We next show the all-or-nothing property for ties. Consider a maximal class B of boxes with a common positive ratio

$$\frac{p_i}{c_i} = \lambda > 0 \quad \text{for all } i \in B.$$

Suppose, toward a contradiction, that a pure best reply opens some but not all boxes in B . Let $k \in B$ be the last opened box from B , and let $\ell \in B$ be an unopened box. By Step 1, no box j with a higher ratio $\frac{p_j}{c_j} > \lambda$ can be opened after k . Similarly, no box m with a lower ratio $\frac{p_m}{c_m} < \lambda$ can be opened after k , since ℓ is left unopened. Hence the rule stops after k if no prize is found.

Let $D > 0$ be the probability of reaching k . Since deleting k cannot improve the rule,

$$\frac{p_k}{D} \cdot \bar{u} - c_k \geq 0 \quad \Longleftrightarrow \quad \lambda \bar{u} \geq D.$$

Since adding ℓ after k cannot improve the rule,

$$\frac{p_\ell}{D - p_k} \cdot \bar{u} - c_\ell \leq 0 \quad \Longleftrightarrow \quad \lambda \bar{u} \leq D - p_k.$$

This is impossible because $p_k > 0$. Therefore, for every positive-ratio class, a pure best reply opens either all boxes in the class or none of them. $\square$

A.3 Proof of Lemma 3

Suppose, toward a contradiction, that

$$\frac{p_i}{c_i} < \frac{p_j}{c_j} \quad \text{for some } i < j.$$

Then $p_j > 0$, so state j lies in Nature's support.

Let b be any pure rule in the support of the DM's mixed strategy. Since $(a, \mathbf{p})$ is a saddle point, every such b is a best reply to $\mathbf{p}$. By Lemma 2, b cannot open i while leaving j unopened, and, if it opens both boxes, it opens j before i . Therefore the payoff of b in state j is weakly larger than its payoff in state i : either b never opens i , in which case state j is weakly better for the DM, or b opens j before i , in which case the prize is found weakly

earlier in state j . Averaging over the DM's mixed strategy gives

$$u(a, j) \geq u(a, i).$$

Next, since $c_i < c_j$, the oracle payoff is strictly larger in state i :

$$u(a^*(i), i) = \bar{u} - c_i > \bar{u} - c_j = u(a^*(j), j).$$

Thus

$$u(a^*(i), i) - u(a, i) > u(a^*(j), j) - u(a, j).$$

Nature assigns positive probability to state j , but moving that probability from state j to state i would strictly increase expected regret against a . This contradicts the optimality of $\mathbf{p}$ at a saddle point. $\square$

A.4 Proof of Theorem 1

Let $(a, \mathbf{p})$ be a saddle point of the reduced game, where a is the DM's mixed strategy. Write

$$\lambda_i := \frac{p_i}{c_i}.$$

By Lemma 3, the sequence $(\lambda_i)_{i=1}^n$ is weakly decreasing.

There is at least one active prize state: if, instead, Nature were to put all probability on state 0, the DM's unique best reply would be to stop immediately, and Nature could obtain strictly positive regret by assigning the prize to box 1, since $c_1 < \bar{u}$. Let

$$m := \max\{i \in [n] : p_i > 0\}.$$

If $p_j > 0$ and $i < j$, then $\lambda_i \geq \lambda_j > 0$, so $p_i > 0$. Hence the active boxes form an initial segment $[m]$. Partition $[m]$ into maximal contiguous intervals on which λ_i is constant, and denote them by

$$B^{(1)}, B^{(2)}, \dots, B^{(K)}.$$

Let $\lambda^{(k)}$ be the common value of p_i/c_i on $B^{(k)}$. Since (λ_i) is weakly decreasing and the blocks are maximal,

$$\lambda^{(1)} > \lambda^{(2)} > \dots > \lambda^{(K)}.$$

Since a is a best reply to $\mathbf{p}$, every pure rule b in its support is itself a best reply to $\mathbf{p}$. By Lemma 2, b opens boxes in weakly decreasing order of p_i/c_i and opens either all boxes or

no boxes from each positive-ratio class. On the active set $[m]$, these positive-ratio classes are precisely the blocks $B^{(1)}, B^{(2)}, \dots, B^{(K)}$. Therefore every pure rule in the support of a searches the active blocks in their order and, for each block, searches either all boxes in the block or none of them. $\square$

A.5 Accounting Identity and Target Vector

Lemma A.1 (Accounting identity). *Fix a nonempty set $S \subseteq [n]$ and an exhaustive order $\pi = (\pi_1, \dots, \pi_m)$ of S . Let*

$$E_{\pi_k}^\pi := \sum_{r=1}^k c_{\pi_r}$$

be the cumulative search cost incurred under this order when the prize is in box π_k . Then

$$\sum_{i \in S} c_i E_i^\pi = \frac{C(S)^2 + \Sigma(S)}{2} = C(S)\Phi(S). \quad (17)$$

Consequently, if the prize is distributed over S according to $\Pr(i \mid i \in S) = c_i/C(S)$, the expected cumulative cost of locating the prize under any exhaustive order is $\Phi(S)$.

Proof. For the order π , we have

$$\sum_{i \in S} c_i E_i^\pi = \sum_{k=1}^m c_{\pi_k} \sum_{r=1}^k c_{\pi_r}.$$

The right-hand side contains each square term c_i^2 once and, for each unordered pair $\{i, j\}$, contains exactly one product $c_i c_j$, namely the product corresponding to whichever of i and j is searched later. Hence

$$\sum_{i \in S} c_i E_i^\pi = \sum_{i \in S} c_i^2 + \sum_{\{i, j\} \subseteq S} c_i c_j = \frac{1}{2} \left[\left(\sum_{i \in S} c_i \right)^2 + \sum_{i \in S} c_i^2 \right] = \frac{C(S)^2 + \Sigma(S)}{2}.$$

Dividing by $C(S)$ gives the expected cumulative cost under the cost-proportional distribution. $\square$

Target vector: Fix block $B^{(k)}$ and consider equation (9), guaranteeing regret equalization between any two boxes contained in $B^{(k)}$. Condition (9) pins down E_i up to an additive constant $a^{(k)}$:

$$E_i = \frac{c_i}{Q^{(k)}} + a^{(k)}, \quad i \in B^{(k)}. \quad (18)$$

The constant $a^{(k)}$ is determined by the previous accounting identity. Namely, any lottery over exhaustive orders of $B^{(k)}$ must satisfy

$$\sum_{i \in B^{(k)}} c_i E_i = C(B^{(k)}) \Phi(B^{(k)}). \quad (19)$$

Substituting (18) into (19) gives

$$a^{(k)} = \Phi(B^{(k)}) - \frac{\Sigma(B^{(k)})}{Q^{(k)} C(B^{(k)})}.$$

Finally, substituting for $a^{(k)}$ yields target vector (10).

A.6 Proof of Theorem 2: Cases and Reach Probabilities

The proof constructs Nature's distribution and the within-tail search lottery and verifies their implementability and optimality. Because these verifications are lengthy and largely algebraic, we relegate them to Online Appendix B.5 but, for later reference, record the case distinction and reach probabilities here. Throughout, write $L_{\ell:n} := \Sigma_{\ell:n}/C_{\ell:n}$.

Case A ($p_0 > 0$): This case applies if $\ell^* = 1$, or if $\ell^* \geq 2$ and $\bar{u} \leq \bar{u}_{\ell^*-1:n}$. The DM randomizes between entering the tail and stopping. Her reach probabilities of the tail and last singleton block are

$$Q^{(\ell^*)} = Q_A^T(\ell^*), \quad \text{and for } \ell^* \geq 2, \quad Q^{(\ell^*-1)} = Q_A^{-T}(\ell^*), \quad (20)$$

where

$$Q_A^T(\ell) := \frac{\bar{u} - L_{\ell:n}}{\bar{u} + C_{\ell:n} - \Phi([\ell, n])}, \quad \text{and for } \ell \geq 2, \quad Q_A^{-T}(\ell) := \frac{\bar{u} - c_{\ell-1} - Q_A^T(\ell) C_{\ell:n}}{\bar{u}}.$$

For $\ell^* \geq 3$, set the remaining reach probabilities recursively

$$Q^{(i)} = \frac{c_{i+1} - c_i + Q^{(i+1)}(\bar{u} - c_{i+1})}{\bar{u}}, \quad i = 1, \dots, \ell^* - 2. \quad (21)$$

Case B ($p_0 = 0$): This case applies if $\ell^* \geq 2$ and $\bar{u} > \bar{u}_{\ell^*-1:n}$. By the minimality of ℓ^* and genericity, these conditions imply $\bar{u} < \Phi([\ell^* - 1, n])$. Conditional on reaching the tail, the DM searches it with probability one. The reach probabilities of the tail and last singleton block are

$$Q^{(\ell^*)} = Q^{(\ell^*-1)} = Q_B^T(\ell^*), \quad (22)$$

where

$$Q_B^T(\ell) := \frac{2(\Sigma_{\ell:n} - c_{\ell-1}C_{\ell:n})}{C_{\ell:n}^2 + \Sigma_{\ell:n}}.$$

For $i = 1, \dots, \ell^* - 2$, set $Q^{(i)}$ as in (21).

A.7 Proof of Theorem 3

Let a and $\tilde{a}$ be the saddle-point search rules selected by Theorem 2 in the original and expanded problems, and let R^* and $\tilde{R}^*$ be the corresponding saddle-point regrets. Let ℓ and $\tilde{\ell}$ denote the associated tail cutoffs. Since the new box is inserted after box 1, the cost of the first box is still c_1 .

First, we argue that $\tilde{R}^* \geq R^*$. By choosing the saddle-point strategy of the original problem, thereby ignoring the newly added box, Nature can guarantee a regret of R^* . The saddle-point regret in the expanded problem must therefore be weakly larger.

Crucially, in both saddle-point constructions, state 1 is active. This implies that the regret in state 1 under the optimal search rule equals the respective saddle-point regret R^* and $\tilde{R}^*$:

$$R^* = R^{(1)} = \bar{u} - c_1 - u(a, 1), \quad \tilde{R}^* = \tilde{R}^{(1)} = \bar{u} - c_1 - u(\tilde{a}, 1).$$

Therefore

$$u(\tilde{a}, 1) = \bar{u} - c_1 - \tilde{R}^* \leq \bar{u} - c_1 - R^* = u(a, 1). \quad (23)$$

Case 1: $\tilde{\ell} > 1$. Then box 1 is the first singleton block in the expanded problem. Hence

$$u(\tilde{a}, 1) = \tilde{Q}^{(1)}(\bar{u} - c_1).$$

In the original problem, conditional on state 1, the DM obtains at most $\bar{u} - c_1$ whenever she initiates search and obtains zero if she does not initiate. Thus

$$u(a, 1) \leq Q^{(1)}(\bar{u} - c_1).$$

Combining this inequality with (23) gives

$$\tilde{Q}^{(1)}(\bar{u} - c_1) \leq Q^{(1)}(\bar{u} - c_1).$$

Since $\bar{u} > c_1$, we obtain $\tilde{Q}^{(1)} \leq Q^{(1)}$.

Case 2: $\tilde{\ell} = 1$. Define the initiation probability that equalizes the regret in state 0 with

single- $[n]$ -block search for the original menu by

$$Q_{[n]} := Q_A^T(1) = \frac{\bar{u} - L_{1:n}}{\bar{u} + C_{1:n} - \Phi([n])}.$$

Let $[\tilde{n}] := [n] \cup \{\text{new}\}$ denote the expanded set. Since $\tilde{\ell} = 1$, the expanded problem is a one-block problem, so

$$\tilde{Q}^{(1)} = Q_{[\tilde{n}]} := \frac{\bar{u} - L_{1:\tilde{n}}}{\bar{u} + C_{1:n} + c_{\text{new}} - \Phi([\tilde{n}])}, \quad \text{with} \quad L_{1:\tilde{n}} = \frac{\Sigma_{1:n} + c_{\text{new}}^2}{C_{1:n} + c_{\text{new}}}.$$

Using $\Phi([n]) = (C_{1:n} + L_{1:n})/2$ and $\Phi([\tilde{n}]) = (C_{1:n} + c_{\text{new}} + L_{1:\tilde{n}})/2$, direct simplification gives

$$Q_{[\tilde{n}]} \leq Q_{[n]} \iff \bar{u} \geq \frac{2C_{1:n}L_{1:n} - c_{\text{new}}(C_{1:n} - L_{1:n})}{C_{1:n} - L_{1:n} + 2c_{\text{new}}}. \quad (24)$$

The condition $\tilde{\ell} = 1$ also implies $\bar{u} \geq \bar{u}_{1:\tilde{n}}$. Another direct simplification gives

$$\bar{u}_{1:\tilde{n}} - \frac{2C_{1:n}L_{1:n} - c_{\text{new}}(C_{1:n} - L_{1:n})}{C_{1:n} - L_{1:n} + 2c_{\text{new}}} = \frac{(c_{\text{new}} - c_1)(C_{1:n}^2 + C_{1:n}L_{1:n} + 2C_{1:n}c_{\text{new}} + 2c_{\text{new}}^2)}{C_{1:n}(C_{1:n} - L_{1:n} + 2c_{\text{new}})} > 0,$$

because $c_{\text{new}} > c_1$. Hence (24) holds, and therefore

$$\tilde{Q}^{(1)} \leq Q_{[n]}. \quad (25)$$

It remains to compare $Q_{[n]}$ with the original $Q^{(1)}$. If the original problem also has $\ell = 1$, then the original construction is the one-block construction, so $Q^{(1)} = Q_{[n]}$. Suppose instead that $\ell > 1$. Since $\tilde{\ell} = 1$, we have $\bar{u} > \Phi([n])$. The function Φ is increasing under inclusion, so $\bar{u} > \Phi([n])$. Consider then the following feasible distribution for Nature in the original problem:

$$\check{p}_0 = \frac{\bar{u} - \Phi([n])}{\bar{u} + C_{1:n} - \Phi([n])}, \quad \check{p}_i = \frac{c_i}{\bar{u} + C_{1:n} - \Phi([n])}, \quad i = 1, \dots, n.$$

All ratios $\check{p}_i/c_i$ are equal. By Lemma 2, the DM either searches the whole block $[n]$ or stops. Stopping gives payoff zero, and exhaustive search also gives payoff zero under $\check{\mathbf{p}}$, since

$$\begin{aligned} & \sum_{i \in [n]} \check{p}_i (\bar{u} - \Phi([n])) - \check{p}_0 C_{1:n} \\ &= \frac{C_{1:n}}{\bar{u} + C_{1:n} - \Phi([n])} (\bar{u} - \Phi([n])) - \frac{\bar{u} - \Phi([n])}{\bar{u} + C_{1:n} - \Phi([n])} C_{1:n} = 0. \end{aligned}$$

Thus Nature can guarantee expected regret

$$\sum_{i=1}^n \check{p}_i(\bar{u} - c_i) = \frac{\bar{u}C_{1:n} - \Sigma_{1:n}}{\bar{u} + C_{1:n} - \Phi([n])} = Q_{[n]}C_{1:n}.$$

Therefore

$$R^* \geq Q_{[n]}C_{1:n}. \quad (26)$$

Because $\bar{u} > \Phi([n])$, the original concave-tail construction is in the positive- p_0 case. Hence state 0 is active and

$$R^* = R_0 = \sum_{i=1}^n c_i \Pr_a(\text{box } i \text{ is opened in state 0}).$$

Since every box is opened with probability at most $Q^{(1)}$, we have

$$R^* \leq Q^{(1)}C_{1:n}.$$

Together with (26), this gives $Q_{[n]} \leq Q^{(1)}$. Combining with (25), we obtain

$$\tilde{Q}^{(1)} = Q_{[\tilde{n}]} \leq Q_{[n]} \leq Q^{(1)},$$

completing the proof. □

A.8 Proof of Theorem 4

For $t \in [0, 1]$, set

$$c_i(t) := c_i + t(\tilde{c}_i - c_i), \quad i = 1, \dots, n,$$

and let $Q^{(1)}(t)$ denote the probability of initiating search in the saddle-point construction of Theorem 2 for the cost profile $c(t)$. Thus

$$Q^{(1)}(0) = Q^{(1)}, \quad Q^{(1)}(1) = \tilde{Q}^{(1)}.$$

Writing $d_i := \tilde{c}_i - c_i$, the assumptions of the theorem imply

$$d_1 \geq d_2 \geq \dots \geq d_n \geq 0,$$

and every $c(t)$ is strictly increasing and concave. We use the following lemma. Its proof is given in Appendix B.6.

Lemma A.2. For $t \in [0, 1]$, let $c_i(t) := c_i + t(\tilde{c}_i - c_i)$. In the following statements, all cost-dependent expressions are evaluated at $c(t)$.

(i) For every ℓ ,

$$\frac{d\Phi([\ell, n])}{dt} \geq 0, \quad \bar{u} \geq \bar{u}_{\ell:n} \implies \frac{dQ_A^T(\ell)}{dt} \leq 0.$$

For $\ell \geq 2$,

$$\bar{u} \geq \Phi([\ell, n]) \implies \frac{dQ_A^{-T}(\ell)}{dt} \leq 0, \quad \frac{dQ_B^T(\ell)}{dt} \leq 0.$$

(ii) Implementability cutoffs are nested:

$$\bar{u}_{\ell:n} \leq \bar{u}_{\ell-1:n}, \quad \ell = 2, \dots, n.$$

(iii) If $\ell \geq 2$ and $\bar{u} = \bar{u}_{\ell-1:n}$, then

$$Q_A^T(\ell) = Q_A^{-T}(\ell) = Q_B^T(\ell).$$

For $\ell = 1, \dots, n-1$, if $\bar{u} = \bar{u}_{\ell:n}$, then

$$Q_A^T(\ell) = Q_B^T(\ell+1).$$

If instead $\bar{u} = \Phi([\ell, n]) \geq \bar{u}_{\ell:n}$, then

$$Q_A^T(\ell) \geq Q_B^T(\ell+1).$$

Finally,

$$Q_B^T(\ell) \geq Q_B^T(\ell+1), \quad \ell = 2, \dots, n-1.$$

Intervals with a fixed construction. Consider an interval on which the tail $[\ell, n]$ and the case of the construction are fixed. In Case A, part (i) of [Lemma A.2](#) shows that $Q_A^T(\ell)$, and also $Q_A^{-T}(\ell)$ when $\ell \geq 2$, are weakly decreasing. In Case B, the same part shows that $Q_B^T(\ell)$ is weakly decreasing. Since

$$Q^{(\ell-1)} = Q^{(\ell)} = Q_B^T(\ell)$$

in Case B, the boundary reach probability is weakly decreasing in either case. Moving backward through the singleton blocks, differentiation of (21) gives

$$\frac{dQ^{(i)}}{dt} = \frac{d_{i+1} - d_i + \frac{dQ^{(i+1)}}{dt}(\bar{u} - c_{i+1}) - Q^{(i+1)}d_{i+1}}{\bar{u}}.$$

Thus $dQ^{(i)}/dt \leq 0$ whenever $dQ^{(i+1)}/dt \leq 0$. Backward induction shows that $Q^{(1)}(t)$ is weakly decreasing on every interval on which the construction is fixed.

Switching points. It remains to show that a change of construction cannot create an upward jump. First, if the case changes while the tail remains $[\ell, n]$, the switch occurs at $\bar{u} = \bar{u}_{\ell-1:n}$. Part (iii) of Lemma A.2 gives $Q_A^T(\ell) = Q_A^{-T}(\ell) = Q_B^T(\ell)$, so the boundary value is continuous.

Second, suppose the tail expands from $[\ell+1, n]$ to $[\ell, n]$. Part (i) of the lemma shows that $\Phi([\ell, n])$ is weakly increasing in t . Hence the larger tail cannot become admissible because its profitability condition becomes newly satisfied. The switch must occur at $\bar{u} = \bar{u}_{\ell:n}$. At this value, part (iii) gives $Q_B^T(\ell+1) = Q_A^T(\ell)$. Moreover, part (ii) implies $\bar{u} \leq \bar{u}_{\ell-1:n}$ when $\ell \geq 2$; for $\ell = 1$, Case A applies by definition. Thus $Q_A^T(\ell)$ is the new reach probability of box ℓ , while $Q_B^T(\ell+1)$ is its old reach probability. There is therefore no upward jump.

Third, suppose the tail contracts from $[\ell, n]$ to $[\ell+1, n]$. If implementability is the binding condition, then $\bar{u} = \bar{u}_{\ell:n}$. Part (ii) implies that the old construction is in Case A when $\ell \geq 2$, while Case A applies by definition when $\ell = 1$. Part (iii) again shows that the old and new reach probabilities of box ℓ agree. If instead profitability is the binding condition, then $\bar{u} = \Phi([\ell, n])$. When implementability binds simultaneously, the preceding argument applies. Otherwise $\bar{u} > \bar{u}_{\ell:n}$, so the construction with the smaller tail $[\ell+1, n]$ is in Case B. If the old construction is in Case A, part (iii) gives $Q_A^T(\ell) \geq Q_B^T(\ell+1)$. If the old construction is in Case B, the final inequality in part (iii) gives $Q_B^T(\ell) \geq Q_B^T(\ell+1)$. Thus a contraction of the tail cannot increase the reach probability of box ℓ .

The recursion (21) is continuous and increasing in $Q^{(i+1)}$. Hence the absence of an upward jump at the tail boundary propagates through every preceding singleton block. Nonadjacent changes can be decomposed into adjacent ones, and simultaneous boundary equalities follow from the same identities by taking one-sided limits. Consequently, $Q^{(1)}(t)$ is weakly decreasing on $[0, 1]$, and therefore

$$\tilde{Q}^{(1)} = Q^{(1)}(1) \leq Q^{(1)}(0) = Q^{(1)}.$$

□

A.9 Proof of Theorem 5

For $k = 1, \dots, n$, write $T_k := [n - k + 1, n]$ for the tail of size k . As $\nu \downarrow 0$,

$$C^\nu(T_k) \rightarrow kc, \quad \Sigma^\nu(T_k) \rightarrow kc^2, \quad \Phi^\nu(T_k) \rightarrow \frac{k+1}{2}c.$$

Moreover,

$$\bar{u}_{n-k+1:n}^\nu \rightarrow c.$$

For $k = 1$, this is immediate from the singleton definition. For $k \geq 2$, it follows from

$$\frac{\Sigma^\nu(T_k)}{C^\nu(T_k)} - c_{n-k+1}^\nu \rightarrow 0$$

while

$$C^\nu(T_k)^2 - \Sigma^\nu(T_k) \rightarrow k(k-1)c^2 > 0.$$

Part (i): If $\bar{u} > \frac{n+1}{2}c$, then, for all sufficiently small ν ,

$$\bar{u} > \Phi^\nu([n]) \quad \text{and} \quad \bar{u} \geq \bar{u}_{1:n}^\nu.$$

Hence the whole set is feasible in (5), so $\ell_\nu = 1$ eventually. The one-block formula in the construction of Theorem 2 gives

$$Q_\nu^{(1)} = \frac{\bar{u} - \Sigma_{1:n}^\nu / C_{1:n}^\nu}{\bar{u} + C_{1:n}^\nu - \Phi^\nu([n])}.$$

Taking limits yields

$$Q_\nu^{(1)} \rightarrow \frac{\bar{u} - c}{\bar{u} + nc - \frac{n+1}{2}c} = \frac{\bar{u} - c}{\bar{u} + \frac{n-1}{2}c}.$$

Part (ii): Now suppose $\bar{u} < \frac{n+1}{2}c$. Then $\bar{n} < n$. Set

$$m := n - \bar{n} + 1.$$

For every tail T_k with $k > \bar{n}$,

$$\frac{k+1}{2}c > \bar{u}.$$

Therefore, for all sufficiently small ν ,

$$\Phi^\nu(T_k) > \bar{u}, \quad k > \bar{n}.$$

No tail of size larger than $\bar{n}$ can then satisfy the profitability condition in (5). Hence $\ell_\nu \geq m$ eventually.

We now use the reach-probability bounds from the construction in [Theorem 2](#). In both Case A and Case B, every singleton block before the tail satisfies

$$0 \leq Q_\nu^{(i)} \leq \frac{c_{i+1}^\nu - c_i^\nu}{c_{i+1}^\nu}, \quad i < \ell_\nu.$$

This is exactly the backward-induction bound used to prove feasibility of the singleton-front recursion. Since $c_{i+1}^\nu - c_i^\nu \rightarrow 0$, the right-hand side converges to zero.

Thus, if $i \leq m$ and $i < \ell_\nu$, then $Q_\nu^{(i)} \rightarrow 0$. It remains only to handle the possible boundary case $i = m = \ell_\nu$. In that case the enlarged tail $[m - 1, n]$ has size $\bar{n} + 1$, and

$$\bar{u}_{m-1:n}^\nu \rightarrow c < \bar{u}.$$

Thus, for all sufficiently small ν , the larger tail fails because it is not profitable: $\bar{u} < \Phi^\nu([m - 1, n])$. The construction is therefore Case B. The concavity bound for the tail-reach probability (22) gives

$$0 \leq Q_\nu^{(m)} \leq \frac{c_m^\nu - c_{m-1}^\nu}{c_m^\nu} \rightarrow 0.$$

Therefore $Q_\nu^{(k)} \rightarrow 0$ for every block index $k = 1, \dots, m = n - \bar{n} + 1$. □

References

- Banchio, M. and Malladi, S. (2025). Rediscovery. *arXiv preprint arXiv:2504.19761*.
- Bao, R., Li, S., and Yu, J. (2023). Search with interdependent values. Working paper.
- Beyhaghi, H. and Cai, L. (2024). Recent developments in pandora’s box problem: Variants and applications. *ACM SIGecom Exchanges*, 21(1):20–34.
- Beyhaghi, H. and Kleinberg, R. (2019). Pandora’s problem with nonobligatory inspection. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 131–132.
- Bizzarri, M. and Lomys, N. (2024). Pandora’s box problem with correlations: Some results for the case of stochastic dominance. Working Paper 740, CSEF – Centre for Studies in Economics and Finance.
- Buturak, G. and Evren, Ö. (2017). Choice overload and asymmetric regret. *Theoretical Economics*, 12(3):1029–1056.

- Chawla, S., Gergatsouli, E., McMahan, J., and Tzamos, C. (2021). Approximating pandora’s box with correlations. *arXiv preprint arXiv:2108.12976*.
- Chawla, S., Gergatsouli, E., Teng, Y., Tzamos, C., and Zhang, R. (2020). Pandora’s box with correlations: Learning and approximation. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1214–1225. IEEE.
- Chernev, A., Böckenholt, U., and Goodman, J. (2015). Choice overload: A conceptual review and meta-analysis. *Journal of Consumer Psychology*, 25(2):333–358.
- Choi, M. and Smith, L. (2020). Optimal sequential search. *Available at SSRN 3530526*.
- Cox, J. C. and Oaxaca, R. L. (1989). Laboratory experiments with a finite-horizon job-search model. *Journal of Risk and Uncertainty*, 2:301–329.
- de Lara, L. and Dean, M. (2024). Rational choice overload.
- Dean, M. et al. (2008). Status quo bias in large and small choice sets. *Unpublished working paper*.
- Dean, M., Ravindran, D., and Stoye, J. (2022). A better test of choice overload. *arXiv preprint arXiv:2212.03931*.
- Deb, J. and Zhou, J. (2018). Reference dependence and choice overload.
- Doval, L. (2018). Whether or not to open pandora’s box. *Journal of Economic Theory*, 175:127–158.
- Edmonds, J. (2003). Submodular functions, matroids, and certain polyhedra. In *Combinatorial Optimization—Eureka, You Shrink! Papers Dedicated to Jack Edmonds 5th International Workshop Aussois, France, March 5–9, 2001 Revised Papers*, pages 11–26. Springer.
- Einav, L. (2005). Informational asymmetries and observational learning in search. *Journal of Risk and Uncertainty*, 30(3):201–226. Working-paper.
- Gambato, J. (2024). Search on a grid: Directed consumer search with correlated products. ZEW Discussion Paper 24-032, ZEW – Leibniz Centre for European Economic Research.
- Gergatsouli, E. and Tzamos, C. (2024). Weitzman’s rule for pandora’s box with correlations. *Advances in Neural Information Processing Systems*, 36.

- Hayashi, T. (2011). Context dependence and consistency in dynamic choice under uncertainty: the case of anticipated regret. *Theory and decision*, 70:399–430.
- Haynes, G. A. (2009). Testing the boundaries of the choice overload phenomenon: The effect of number of options and time pressure on decision difficulty and satisfaction. *Psychology & Marketing*, 26(3):204–212.
- Ichiishi, T. (1981). Super-modularity: applications to convex games and to the greedy algorithm for lp. *Journal of Economic Theory*, 25(2):283–286.
- Inbar, Y., Botti, S., and Hanks, K. (2011). Decision speed and choice regret: When haste feels like waste. *Journal of Experimental Social Psychology*, 47(3):533–540.
- Iyengar, S. S. and Kamenica, E. (2010). Choice proliferation, simplicity seeking, and asset allocation. *Journal of Public Economics*, 94(7-8):530–539.
- Iyengar, S. S. and Lepper, M. R. (2000). When choice is demotivating: Can one desire too much of a good thing? *Journal of personality and social psychology*, 79(6):995.
- Kamenica, E. (2008). Contextual inference in markets: On the informational content of product lines. *American Economic Review*, 98(5):2127–2149.
- Ke, T. T. and Lin, S. (2020). Informational complementarity. *Management Science*, 66(8):3699–3716.
- Keller, G. and Oldale, A. (2003). Branching bandits: A sequential search process with correlated pay-offs. *Journal of Economic Theory*, 113(2):302–315.
- Kogut, C. A. (1990). Consumer search behavior and sunk costs. *Journal of Economic Behavior & Organization*, 14(3):381–392.
- Long, X., Sun, J., Dai, H., Zhang, D., Zhang, J., Chen, Y., Hu, H., and Zhao, B. (2025). The choice overload effect in online recommender systems. *Manufacturing & Service Operations Management*, 27(1):249–268.
- Malladi, S. (2025). Searching in the dark and learning where to look. *Available at SSRN 4084113*.
- Nocke, V. and Rey, P. (2024). Consumer search, steering, and choice overload. *Journal of Political Economy*, 132(5):1684–1739.
- Olszewski, W. and Weber, R. (2015). A more general pandora rule? *Journal of Economic Theory*, 160:429–437.

- Reutskaja, E. and Hogarth, R. M. (2009). Satisfaction in choice as a function of the number of alternatives: When “goods satiate”. *Psychology & Marketing*, 26(3):197–203.
- Rothschild, M. (1974). A two-armed bandit theory of market pricing. *Journal of Economic Theory*, 9(2):185–202.
- Sarver, T. (2008). Anticipating regret: Why fewer options may be better. *Econometrica*, 76(2):263–305.
- Scheibehenne, B., Greifeneder, R., and Todd, P. M. (2010). Can there ever be too many options? a meta-analytic review of choice overload. *Journal of consumer research*, 37(3):409–425.
- Schlag, K. H. and Zapechelnyuk, A. (2021). Robust sequential search. *Theoretical Economics*, 16(4):1431–1470.
- Schotter, A. and Braunstein, Y. M. (1981). Economic search: An experimental study. *Economic Inquiry*, 19(1):1–25.
- Sethi-Iyengar, S., Huberman, G., Jiang, W., et al. (2004). How much choice is too much? contributions to 401(k) retirement plans. *Pension design and structure: New lessons from behavioral finance*, 83:84–87.
- Shah, A. M. and Wolford, G. (2007). Buying behavior as a function of parametric variation of number of choices. *Psychological Science*, 18(5):369–370.
- Shapley, L. S. (1971). Cores of convex games. *International journal of game theory*, 1(1):11–26.
- Sonnemans, J. (1998). Strategies of search. *Journal of Economic Behavior & Organization*, 35(3):309–332.
- Weitzman, M. L. (1979). Optimal search for the best alternative. *Econometrica*, 47(3):641–654.