Moonshots, Monkeys, and Pedestals: Optimal Stopping and Prioritization in Multi-Component Projects

Preliminary version; June 29 2026. Please do not distribute nor cite without permission.

Vladimir Smirnov
University of Sydney, School of Economics, NSW 2006 Australia
vladimir.smirnov@sydney.edu.au

Ilia Tsetlin
INSEAD, 1 Ayer Rajah Avenue, Singapore 138676
ilia.tsetlin@insead.edu

Saša Zorc
University of Virginia, Darden School of Business, 100 Darden Blvd, Charlottesville, VA 22903
zorcs@darden.virginia.edu

We investigate the optimal strategy to execute multi-component projects, where completing all components is necessary to obtain a reward. Each component might be impossible to complete. This setting is typical for various R&D activities, including software development and academic research. We study a resulting dynamic optimal stopping problem, that involves the decision which component to work on and when to stop. If each component is described by prior probability that it is feasible and constant execution rate, the optimal strategy is to work on one component and abandon the whole project if this component is not completed by the optimal stopping time. We also explore which component should be the one to start with, deriving an approximate index policy and showing that for moonshot projects one should start with the component with the highest execution rate.

1. Introduction

Innovation projects are often high-upside, high-uncertainty endeavours in which value is realised only if several necessary components are successfully completed. A drug-development project may require both a scientific breakthrough and regulatory approval; a software product may require both a working core technology and a scalable user interface; an academic project may require both a tractable theoretical mechanism and suitable data. In such settings, each component may ultimately prove infeasible. This creates two closely related managerial questions. Which component should be attempted first? And how long should the decision-maker continue working before abandoning the project?

Consider the proverbial "monkey and the pedestal" example, originally introduced by Google X's Eric 'Astro' Teller (Teller 2016), and widely popularized in publications for practicing managers (e.g., Duke 2022, Barker 2023, Kahl 2023). In this example, you are tasked with building an ornate pedestal on which stands a monkey that recites Shakespeare. The pedestal component appears relatively straightforward, while training the monkey to recite Shakespeare is far more complex and may even be impossible. So, what should one work on first? A natural temptation is to start with the visible and more tractable part of the project: the pedestal. However, as argued in Teller (2016), this may be exactly the wrong approach. If one builds the pedestal first, but then fails to train the monkey, all effort spent on the pedestal has been wasted. The recommendation is therefore to do the hard and uncertain part first, leading to the hashtag **#monkeyfirst**.

This logic is compelling, but it is not complete. If the pedestal is known to be feasible, then it can indeed be safely postponed. But if completing the pedestal is also uncertain, even slightly, should one ever start with it? If the pedestal can be tested very quickly or cheaply, it may be worth resolving its uncertainty before spending resources on the harder component. More generally, the relevant question is not simply whether a component is hard or easy, or whether it is likely or unlikely to succeed. The optimal order depends on the interaction between feasibility, speed of completion, cost of experimentation, and the option to abandon the project.

This paper connects two literatures that are usually studied separately. The optimal-stopping and project-abandonment literature studies when unsuccessful experimentation should lead a decision-maker to stop. The sequencing literature studies which tasks, experiments, or project components should be undertaken first. Our setting combines the two problems: because all components are necessary for success, the decision-maker must jointly decide which component to test and when continued lack of success makes the entire project unattractive. Section 2 discusses these literatures in more detail.

Returning to the monkey and the pedestal example, suppose that the cost of training the monkey is c_M, the cost of building the pedestal is c_P, and the probabilities of success are p_M and p_P, respectively. If each component can be tried only once, then attempting the monkey first is better than attempting the pedestal first if

$$c_M + p_M c_P < c_P + p_P c_M,$$

or equivalently,

$$\frac{1-p_M}{c_M} > \frac{1-p_P}{c_P}.$$

Thus, in the one-shot version of the problem, the decision-maker should first attempt the component with the higher probability of failure per unit cost. This captures the basic "test the weak link first" logic behind the monkey-first recommendation. However, if components can be explored over time, the problem changes. Lack of success is not an immediate failure, but it is informative. It gradually changes beliefs about feasibility. It is then no longer obvious which component should be attempted first, how long it should be attempted, or whether it can ever be optimal to switch between components before abandoning the whole project.

Our paper fills this gap by studying sequencing and abandonment jointly when each necessary component can be explored over time and unsuccessful exploration gradually changes beliefs about feasibility. We study a multi-component project in which all components must be completed before any reward is obtained. Each component is a Poisson bandit: while the decision-maker works on a component, completion arrives according to an unknown arrival rate. The absence of completion is informative, making the decision-maker progressively more pessimistic about that component and, because all components are necessary, about the project as a whole. Although we also consider general prior distributions over success rates, our main analysis focuses on a baseline feasibility-risk model. In this setting, component i is feasible with probability p_i; if feasible, completion arrives according to a Poisson process with known rate λ_i, while if infeasible, completion never arrives. This is the single-component learning structure studied by McCardle et al. (2018).

Our first result shows that the dynamic allocation problem is simpler than it may initially appear. In principle, the decision-maker could repeatedly switch between components, or divide effort across them over time. We show that, under general priors, the problem can be reduced to choosing how long to work on each component.

Our second result is sharper for the feasibility-risk model. In this case, the optimal policy is concentrated on one component. The decision maker starts with a component and, absent success, either continues working on the same component or abandons the entire project. She does not switch to another component after a period of unsuccessful work. The intuition is that lack of success makes the current component look less promising. Since all components are necessary, this also makes the whole project less promising.

Our third result concerns which component should be attempted first. The monkey-first recommendation is sometimes correct, but not always. A component with a lower probability of feasibility is more attractive to test first because it is more likely to reveal that the project should be abandoned. However, the probability of feasibility is not the only relevant consideration. The optimal order also depends on the rate at which completion arrives if the component is feasible. As a result, it can be optimal to begin with the pedestal rather than the monkey, provided that the pedestal is not perfectly certain and can be tested sufficiently quickly.

This reversal is especially important for moonshot projects, where the reward from success is large relative to the costs of experimentation. For sufficiently large rewards, the optimal priority rule becomes lexicographic. The decision maker first chooses the component with the higher arrival rate conditional on feasibility. If arrival rates are equal, she chooses the component with the lower prior probability of feasibility. Hence, in high-reward projects, the apparently easier component may optimally be attempted first because it resolves uncertainty quickly.

We also develop a simple approximate index policy. The index compares the probability of failing to complete a component with the expected cost of testing it. It is easy to compute because it depends only on the primitives of the component being ranked. It also has a natural interpretation: it measures how effectively work on a component can reveal that the entire project should be abandoned. We show that this index performs particularly well for moonshot projects and becomes asymptotically optimal as the reward from success grows.

The rest of the paper proceeds as follows. Section 2 discusses the related literature in more detail. Section 3 introduces the model, solves the single-component stopping problem, and characterises the multi-component allocation problem. It also derives the corner-solution result for the baseline feasibility-risk model and studies the basic prioritisation problem. Section 4 focuses on moonshot projects and characterises the limiting priority rule when the reward from success is large. Section 5 develops the approximate index policy and studies its performance. Section 6 extends the analysis to general prior distributions over success rates. Section 7 gives an overview of what in the paper extends seamlessly to projects with many components and what does not. Section 8 concludes.

2. Related Literature

We now discuss the two literatures most closely related to the paper: optimal stopping and project abandonment under learning, and the sequencing of tasks, experiments, and project components. The distinction is useful because our model combines these problems in an all-or-nothing project environment.

2.1. Optimal stopping, project review, and abandonment under learning

The first related literature studies when a decision-maker should continue, abandon, or revise an uncertain project as information arrives. Roberts and Weitzman (1981) are an early and influential contribution, emphasizing that research and development projects should be evaluated dynamically because experimentation produces information and continuation decisions are revised as this information arrives. McCardle et al. (2018) study when a decision-maker should abandon a research project and search for a new one. Their model captures the trade-off between persistence and abandonment when lack of success makes the decision-maker increasingly pessimistic about eventual feasibility or profitability. Their single-project stopping problem is an important building block for our analysis: once all but one component have been completed, the remaining problem has the same logic of experimentation, belief updating, and abandonment. Thijssen and Bregantini (2017) also study costly sequential experimentation and project valuation when experimentation reveals information about project profitability.

A related set of papers studies project-specific investment, learning, and continuation decisions. Kwon and Lippman (2011) analyze the acquisition of project-specific assets under Bayesian updating, while Kwon (2010) studies optimal invest-or-exit decisions in the face of a declining profit stream. Alizamir et al. (2022) study search under accumulated pressure, combining Bayesian inference, stopping, and dynamic pressure from unresolved tasks. A broader optimal-experimentation literature provides the learning foundation for these problems: Moscarini and Smith (2001) analyze the trade-off between current payoffs and the value of information, and Keller et al. (2005) and Keller and Rady (2010) study continuous-time strategic experimentation with exponential or Poisson bandits.

Recent work on project management, deadlines, and dynamic incentives is also relevant. Green and Taylor (2016) analyze contracting for multistage projects when an agent privately observes intermediate progress. Madsen (2022) studies deadline design for projects of uncertain scope, while Feng et al. (2024) examine project management with repeated setbacks, shutdown decisions, and overruns. Curello and Sinander (2025) study mechanisms for screening breakthroughs when the time of a breakthrough is privately observed. These papers share the view that continuation and termination are dynamic decisions, although their focus is primarily on contracting, private information, reporting, or agency.

The closest paper in this group is Moroni (2022), who studies experimentation inside organizations when a project consists of tasks or milestones of unknown feasibility and beliefs decline after unsuccessful work. This is conceptually close to our feasibility-learning environment. The key difference is that Moroni studies a dynamic moral-hazard problem without sequencing decision, while we study the normative decision problem of a single decision-maker who chooses which component to pursue and when to abandon the project. This allows us to characterize the sequencing and stopping policy directly and to derive the corner-solution result under feasibility uncertainty.

Other recent work positions the paper within the broader economics of innovation and project governance. Kirpalani and Madsen (2023) study strategic investment evaluation when firms privately acquire information and choose when to invest. Guo (2016) and Escobar and Zhang (2021) study dynamic delegation and delegated learning. Carnehl et al. (2025) examine the design of scientific grants and the trade-off between funding duration, risk taking, and research uncertainty. Goldstein and Kearney (2020) provide empirical evidence on active project management in public research funding, documenting how ARPA-E modifies timelines and budgets in response to project performance.

2.2. Sequencing tasks, components, experiments, and index rules

The second related literature studies the optimal order in which tasks, tests, or project components should be undertaken. Agastya and Birulin (2023) is one of the closest economics papers in this category. They study a project consisting of several tasks that differ in success probabilities and incremental benefits and must be assigned to specialists subject to adverse selection and moral hazard. They show that, despite informational constraints, a simple index determines the optimal schedule and mechanism. Their paper is a natural comparator because it studies which task should be performed first and when a project should be terminated. Our model differs by replacing one-shot task outcomes with continuous-time Bayesian learning: unsuccessful work does not immediately reveal failure, but gradually lowers the posterior probability that the component is feasible.

Operations-management work on R&D and project scheduling with activity failures is also directly relevant. De Reyck and Leus (2008) study R&D project scheduling when activities may fail and when failure can affect the value of the overall project. Zhao et al. (2020) analyze project evaluation and selection when the failure of any task can terminate the project. Creemers et al. (2015) consider project planning with alternative technologies in uncertain environments. Schmidt and Grossmann (1996) develop optimization models for scheduling testing tasks in new product development, Erat and Kavadias (2008) analyze sequential testing of product designs, and Thomke and Bell (2001) study the timing, frequency, and fidelity of sequential testing in product development. These papers share our concern with the order in which uncertain project activities should be undertaken; our focus is on the interaction between sequencing and optimal stopping when failure is learned gradually through the absence of completion.

A particularly relevant contribution in this project-task tradition is Kwon et al. (2011), who study sourcing decisions for project tasks with exponential completion times. Their exponential completion times are close to the Poisson-completion primitives in our model. The main distinction is that their focus is on sourcing and task allocation, whereas our focus is on abandonment and prioritization under feasibility uncertainty.

Search and bandit models provide a second foundation for the sequencing part of the paper. Weitzman (1979) introduced the Pandora's box problem, in which a decision-maker chooses the order in which to inspect costly alternatives. Gittins (1979) developed dynamic allocation indices for bandit processes, and Beyhaghi and Cai (2024) survey recent developments in Pandora's box problems and related index methods. These models show when complex dynamic allocation problems can be represented or approximated by simple indices. In our setting, however, components are complements rather than substitutes: all must be completed for the reward to be realized. Thus, the relevant index is not designed to identify the best alternative, but to identify failure quickly and cheaply in an all-or-nothing project.

The sequential-exploration literature is also related. Bickel and Smith (2006) study optimal sequential exploration in a binary learning model, Bickel et al. (2008) analyze dependence among geologic risks in sequential exploration decisions, and Brown and Smith (2013) study sequential exploration problems in which the decision-maker chooses where to explore first and next. Our approximate index policy is also related to Brown and Smith (2020), who study index policies and performance bounds for dynamic selection problems.

Finally, our paper speaks to the literature on innovation search and breakthrough discovery. Fleming (2001) emphasizes that technological search is especially uncertain when innovation requires novel combinations of components. Chaimanowong et al. (2023) study endogenous choice and stopping in a model of searching for breakthroughs across projects. Oliveros (2025) develops a dynamic model of active product development in which a developer can continue, restart, or terminate development of a product of uncertain quality. Carnehl and Schneider (2023) study the allocation of time between acting on a risky idea and searching for alternatives under time pressure, Carnehl and Schneider (2025) analyze how knowledge evolves through the choice of research questions, and Urgun and Yariv (2025) study contiguous search and ambition on uncharted terrain.

Taken together, these literatures highlight the gap filled by the present paper. The project-abandonment literature explains how pessimism generated by unsuccessful experimentation can trigger stopping. The sequencing literature explains how task order matters when tasks differ in cost, difficulty, informativeness, or contribution to value. Our paper combines these ideas in a tractable model of complementary Poisson components with Bayesian learning about feasibility. Under feasibility uncertainty, the decision-maker works on one component until either it succeeds or the project is abandoned. The model also clarifies when the standard managerial recommendation to work on the hardest or most uncertain component first is correct and when it reverses.

3. Model

3.1. Setup

The project consists of n components, all of which must be successfully completed for the decision-maker (DM) to receive the reward R. When working on a component, success arrives via a Poisson process with unknown arrival rate. Formally, each component $i \in \{1,\dots,n\}$ has an unknown arrival rate $\Lambda_i \geq 0$. We allow $\Lambda_i = 0$, in which case the component is *infeasible*: no amount of work will complete it. Working on a component costs c per unit of time, and the random variables $\{\Lambda_i\}_{i=1}^n$ are independent across components. We make the mild assumption $\mathbb{E}[\Lambda_i] < \infty$ for every i, which is needed only to ensure that the value functions below are well-defined.

The two-point "feasibility" model of McCardle et al. (2018) is the special case in which Λ_i equals λ_i with probability p_i and 0 with probability $1 - p_i$. We will recover this special case for specific results in the paper, but generally we allow Λ_i to follow an arbitrary non-negative distribution—for instance, a continuous distribution reflecting uncertainty about how *difficult* a component is, or a mixed distribution combining feasibility uncertainty with rate uncertainty.

The DM updates her beliefs about Λ_i as work proceeds without success. Rather than tracking the full posterior distribution, we encode the Bayesian content of the model through two scalar functions of the cumulative time t spent on component i:

$$\pi_i(t) := \mathbb{E}\big[e^{-\Lambda_i t}\big], \tag{1}$$
$$\hat{\Lambda}_i(t) := \frac{\mathbb{E}[\Lambda_i e^{-\Lambda_i t}]}{\pi_i(t)} = -\frac{d\log\pi_i(t)}{dt}. \tag{2}$$

Both quantities are taken under the prior; no posterior distribution needs to be written down.

The function $\pi_i(t)$ is the prior probability that no success has arrived after t units of *continuous* work on component i, obtained by averaging the conditional survival probability $e^{-\Lambda_i t}$ over the prior. The function $\hat{\Lambda}_i(t)$ is the DM's current best estimate of the arrival rate after t units of work without success: by Bayes' rule it equals the posterior expectation of Λ_i given the event "no success in the first t units of work," and the second equality in (2) expresses this estimate as a derivative of the simpler function π_i. Both $\pi_i(t)$ and $\hat{\Lambda}_i(t)$ depend only on cumulative time spent on component i.

The DM is risk-neutral and discounts future cash flows at rate $\delta > 0$. She faces two intertwined decisions, made in continuous time. *First*, a task-prioritization problem: at every instant she must decide which component to work on, and in principle she may split her effort across components at any rate. *Second*, an optimal-stopping problem: she may terminate the entire project at any time if her posteriors have become pessimistic enough that continuing is no longer worthwhile. The remainder of the paper studies the joint solution of these two decisions.

3.2. A single-component model

Consider first the situation where there is only a single component to be completed, or equivalently, the situation where the DM has successfully completed all but one component. This situation is considerably simpler, as it eliminates the task prioritization component of the model, leaving only the optimal stopping decision. The problem that the DM faces is then analogous to the model of McCardle et al. (2018) so we leverage their solution here. Let T_i denote the time spent on component i, and Λ_i denote the unknown arrival rate of success (Λ_i is a random variable that is equal to λ_i with probability p_i, and it 0 otherwise). The expected value from working on project i for time T_i is given by:

$$V_i(T_i) = \mathbb{E}\left[\int_0^{T_i} \Lambda_i e^{-\Lambda_i t}\left(Re^{-\delta t} - \frac{c}{\delta}(1-e^{-\delta t})\right)dt - \frac{c}{\delta}(1-e^{-\delta T_i})e^{-\Lambda_i T_i}\right]. \tag{3}$$

Solving the integral and simplifying yields the closed-form expression

$$V_i(T_i) = \mathbb{E}\left[\left(\frac{\Lambda_i}{\Lambda_i+\delta}R - \frac{c}{\Lambda_i+\delta}\right)\left(1-e^{-(\Lambda_i+\delta)T_i}\right)\right]. \tag{4}$$

The first-order condition for T_i takes a particularly clean form once we recognize the two functions $\pi_i(\cdot)$ and $\hat{\Lambda}_i(\cdot)$ from (1)–(2). Differentiating (4) and factoring,

$$V_i'(T_i) = \mathbb{E}\big[(\Lambda_i R - c)\, e^{-(\Lambda_i+\delta)T_i}\big] = e^{-\delta T_i}\, \pi_i(T_i)\, \big(R\,\hat{\Lambda}_i(T_i) - c\big). \tag{5}$$

The bracket on the right is the gap between the posterior expected reward rate $R\,\hat{\Lambda}_i(T_i)$ and the cost rate c after T_i time of no success. Setting $V_i' = 0$ therefore yields a clean abandonment rule.

Proposition 1 (Single-component optimum) *The optimal stopping time T_i^* for the single-component problem is as follows:* (i) *If $R\,\mathbb{E}[\Lambda_i] \le c$, then $T_i^* = 0$ (abandon immediately).* (ii) *If $R \inf \operatorname{supp}(\Lambda_i) \ge c$, then $T_i^* = \infty$ (never abandon).* (iii) *Otherwise, T_i^* is the unique time at which*

$$R\,\hat{\Lambda}_i(T_i^*) \;=\; c. \tag{6}$$

We will use this solution as a building block of the full model, as the DM will be in this situation once she has completed all but one component. We denote the single-component value function under the optimal stopping time as $V_i^* := V_i(T_i^*)$.

Lastly we note that while a closed form for T_i^* is not available in general, the two-point case where $\Lambda_i \in \{0, \lambda_i\}$ with $\mathbb{P}(\Lambda_i = \lambda_i) = p_i$ yields the closed form for T_i^*: it is 0 if $p_i R \lambda_i \le c$, and it is otherwise given by

$$T_i^* = -\frac{1}{\lambda_i} \ln\left[\frac{1-p_i}{p_i}\,\frac{c}{R\lambda_i - c}\right]. \tag{7}$$

3.3. A multi-component model

We now proceed to the full multi-component model. While almost all of our results extend to an arbitrary n in a straightforward way (see Section 7 for details); for expositional clarity we focus the paper on $n = 2$. Recall from Section 3.1 that the arrival rates Λ_1 and Λ_2 are independent and that for each component the no-success probability $\pi_i(t)$ and the posterior expected rate $\hat{\Lambda}_i(t)$ are functions of cumulative time spent on that component.

Until a component is completed, the state of the system at time t can be fully described by $x(t) = (t_1(t), t_2(t))$: the time the DM has spent working on each component so far. The decision about which component to work on (or how to spread your time between them) is modeled by the control function $u(t) \in \mathcal{U}$ where $\mathcal{U}$ is the set of all right-differentiable functions from $\mathbb{R}_+$ to $[0,1]$. The function $u(t)$ gives the fraction of time to dedicate to component 1, at time t. The state then starts at $x = (0,0)$ at time 0 and its evolution function is $\dot{x} = (u(t), 1 - u(t))$, thus $x(t) = (\int_0^t u(v)dv, \int_0^t (1 - u(v))dv)$.

After success on component i, the decision-maker faces the single-component problem of Section 3.2 on component $j = 3 - i$ at her current state of beliefs about Λ_j, with continuation value under single-component optimal policy and updated beliefs denoted $V_j^*(t_j)$.

The DM's central problem can then be expressed as optimization over the control function and stopping time as follows

$$\sup_{u\in\mathcal{U},T\in\mathbb{R}_+} \int_0^T g(x,u,t)dt - C(u,T) \tag{8}$$

$$\text{s. t.} \quad \dot{x} = (u(t), 1-u(t)), \tag{9}$$

where $g(x,u,t)$ is the instantaneous reward function and $C(u,T)$ the expected cost function, which are given by

$$g(x,u,t) = e^{-\delta t}\,\pi_1(t_1)\,\pi_2(t_2)\left[u(t)\hat{\Lambda}_1(t_1)\,V_2^*(t_2) + (1-u(t))\hat{\Lambda}_2(t_2)\,V_1^*(t_1)\right], \tag{10}$$

$$C(u,T) = \int_0^T f(x,u,t)c\frac{1-e^{-\delta t}}{\delta}\,dt + \pi_1(t_1(T))\,\pi_2(t_2(T))c\frac{1-e^{-\delta T}}{\delta}, \tag{11}$$

where $f(x,u,t)$ is the density of the first success at time t under the control u, given by

$$f(x,u,t) \;=\; \left[u(t)\,\hat{\Lambda}_1(t_1(t)) + (1-u(t))\,\hat{\Lambda}_2(t_2(t))\right]\pi_1(t_1(t))\,\pi_2(t_2(t)). \tag{12}$$

We elaborate on the source of these expressions next. From the DM's perspective, given her decision u, she faces a non-homogeneous Poisson process via which two types of success can arrive (one for each component). The source of this non-homogeneity is two-fold: the arrival rate changes as the decision-maker alters the allocation of her efforts between the components, and it also changes as her beliefs about each component's Λ_i are updated by the no-success evidence accumulated on that component. The (time-varying) arrival rate of this process is the Bayes-updated effective rate

$$u(t)\,\hat{\Lambda}_1(t_1(t)) + (1-u(t))\,\hat{\Lambda}_2(t_2(t)).$$

Because $\hat{\Lambda}_i = -(\log\pi_i)'$ by (2), a change of variable gives

$$\int_0^t \left[u(v)\,\hat{\Lambda}_1(t_1(v)) + (1-u(v))\,\hat{\Lambda}_2(t_2(v))\right]dv \;=\; -\log\pi_1(t_1(t)) - \log\pi_2(t_2(t)),$$

so the Poisson survival factor $\exp\{-\int_0^t \lambda(v)\,dv\}$ for this effective rate equals $\pi_1(t_1(t))\,\pi_2(t_2(t))$. Applying the standard property that for any non-homogeneous Poisson process with arrival rate $\lambda(t)$ the density of first arrival is $\lambda(t)\,e^{-\int_0^t \lambda(v)dv}$ then yields (12). The instantaneous reward function (10) consists of this density, the factor by which a reward at time t is discounted ($e^{-\delta t}$), and the expected payoff in case of arrival (the expression in brackets in (10), which results from success on component 1 yielding the DM an expected continuation payoff of $V_2^*(t_2(t))$, whereas success on component 2 yields $V_1^*(t_1(t))$). Lastly, the expected cost function (11) is based on the discounted costs of a DM who works until time t using control u and experiences no arrivals being $\int_0^t e^{-\delta v}c\,dv = c(1-e^{-\delta t})/\delta$. The first additive term of (11) accounts for situations where an arrival happens before T (so that the costs experienced depend on the exact time of first arrival, which has density $f(x,u,t)$), whereas the second term accounts for situations where no arrival happens until T, which occurs with probability $\pi_1(t_1(T))\,\pi_2(t_2(T))$.

The theorem below gives our first main result, and shows a control space collapse in (8)–(9) that considerably simplifies the problem.

Theorem 1 *Consider a stopping time $T \in \mathbb{R}_+$ and any two control functions $u, \hat{u} \in \mathcal{U}$ satisfying $\int_0^T u(s)\,ds = \int_0^T \hat{u}(s)\,ds \le T_1^*$ and $T - \int_0^T u(s)\,ds \le T_2^*$. Then, (u,T) and $(\hat{u},T)$ give the same expected value in problem (8)–(9). Any optimal policy $u(t)$ satisfies $\int_0^T u(s)\,ds \le T_1^*$ and $T - \int_0^T u(s)\,ds \le T_2^*$; thus, without loss of optimality, (8)–(9) can be reduced to optimization over $(T_1,T_2) := (\int_0^T u(s)ds, T - \int_0^T u(s)ds)$.*

Essentially, the theorem implies that, for all practical purposes, the problem of distributing effort between the components, or in which order to prioritize the components, does not matter, as they all yield the same expectation.

The key decision that remains is how much to work on each component (T_i) before giving up on the project if none of them yield success. This statement is subject to mild conditions: The $\int_0^T u(s)ds \le T_1^*$ and $T - \int_0^T u(s)ds \le T_2^*$ require that the DM never works on a component for more time than would be optimal if that component was the only component. Notice that any optimal (u,T) pair (one that solves (8)–(9)) has to satisfy this condition. So, while there exist policies under which the best order of execution may be different, none of such policies can be optimal.

This result gives a clear simplification: conditional on the cumulative experimentation times T_1 and T_2, the precise path by which these times are accumulated is payoff-equivalent. The remaining decision is therefore how much cumulative experimentation time to allocate to each component before abandoning the project in the absence of success.

With this result, we can write the DM's value function as under any policy as $V(T_1,T_2)$, which is then given by the symmetric expression (detailed derivation is in Appendix EC.1):

$$\begin{aligned} V(T_1,T_2) = {} & \mathbb{E}_{\Lambda_1}\left[\left(\frac{\Lambda_1}{\Lambda_1+\delta}V_2^* - \frac{c}{\Lambda_1+\delta}\right)\left(1-e^{-(\Lambda_1+\delta)T_1}\right)\right] \\ & + \mathbb{E}_{\Lambda_2}\left[\left(\frac{\Lambda_2}{\Lambda_2+\delta}V_1^* - \frac{c}{\Lambda_2+\delta}\right)\left(1-e^{-(\Lambda_2+\delta)T_2}\right)\right] \\ & - \mathbb{E}_{\Lambda_1,\Lambda_2}\left[\frac{\Lambda_1\Lambda_2 R - c\Lambda_2 - c\Lambda_1 - c\delta}{(\Lambda_1+\delta)(\Lambda_2+\delta)}\left(1-e^{-(\Lambda_1+\delta)T_1}\right)\left(1-e^{-(\Lambda_2+\delta)T_2}\right)\right]. \end{aligned} \tag{13}$$

The problem (8)-(9) thus reduces to

$$\max_{T_1\in[0,T_1^*],\,T_2\in[0,T_2^*]} V(T_1,T_2) \tag{14}$$

Now, this problem is still not trivial as (13) is neither convex nor concave. However, the problem can be simplified further by the following theorem, which establishes that under a relatively weak condition, at any interior stationary point the Hessian of the objective in (14) is indefinite; thus, all interior stationary points are saddles.

Theorem 2 *If the ratio $V_i^*(t_i)/\hat{\Lambda}_i(t_i)$ is decreasing in t_i, (14) cannot have an interior solution.*

The key insight from this theorem is that when the value function is elastic with respect to $\hat{\Lambda}_i$, allocating time to one component reduces the perceived value of continuing the other component more significantly; thus, discouraging an interior solution.

Note that the condition of Theorem 2 is sufficient but not necessary. In particular, it needs to hold only at the critical point where $V_i^*(t_i) = c/\hat{\Lambda}_j(t_j)$ to rule out an interior optimum (see proof for details).

While it may be difficult to intuit when this condition will hold, we can establish that it does universally hold in the canonical two-point case where the uncertainty is not about the speed at which the components can be completed, but about their feasibility. We formalize this in the following result.

Proposition 2 *If* $\text{supp}(\Lambda_i) = \{0, \lambda_i\}$*, the condition of Theorem 2 holds, and thus (14) can only have corner solutions.*

With this in mind, we focus now on the more tractable setup of the two-point case where $\text{supp}(\Lambda_i) = \{0, \lambda_i\}$, and as in (7), denote $p_i := \mathbb{P}(\Lambda_i = \lambda_i)$, and assume this distribution going forth. We will revisit the case of a general distribution for Λ_i in Section 6, where we will show that despite the tractability issues, most of the paper's central results will extend to the general setting.

It is thus optimal for the DM to focus entirely on a single component until either that component is successfully completed or the DM becomes so pessimistic from lack of success that it becomes optimal to terminate the entire project. This finding is very consistent with the approach of "fail fast, fail often, fail cheap" that is often promoted by the tech industry (Sitkin 1992, McGrath 2011, Babineaux and Krumboltz 2013). Or in other words, if the project is doomed to be terminated (due to likely infeasibility of some components) the DM would like to arrive to the failure state in the fastest (cheapest) way possible. Thus, the DM works on the component that looks the *least* promising. As the DM works more on it without success, the component starts looking even less promising due to Bayesian updating; thus, in the absence of success, it never changes which component is the least promising one.

A natural question that arises then is what exactly makes a component the "least promising one?" What is the component the DM should focus on?

With the result of Proposition 2 in mind, the problem reduces to one bearing a strong resemblance to the monkey and the pedestal prioritization question discussed in the introduction. One can think of the proverbial monkey as a task with low λ_i and p_i, while the pedestal is a task for which those parameters are both high. (Of course, the question of prioritization is more difficult if one task has higher p_i but lower λ_i.)

To answer this question, notice that the optimal policy for a DM who focuses on just component i can be easily derived by recursively applying the solution to the single component problem. The situation faced by such DM is equivalent to a single-component model in which the reward for successfully completing the first component is V_j^* instead of R (here j denotes the other component, that is, $j = 3 - i$). Thus, replacing the R in (4), inserting the single-component solution (7) and expanding the resulting expression yields that the value function of such a DM is

$$V_i^{**} = \frac{1}{\lambda_i + \delta}\left(p_i(\lambda_i V_j^* - c) - c(1-p_i)\left(\frac{c(1-p_i)}{p_i(V_j^*\lambda_i - c)}\right)^{\delta/\lambda_i}\right) + \frac{(1-p_i)c}{\delta}\left(\left(\frac{c(1-p_i)}{p_i(V_j^*\lambda_i - c)}\right)^{\delta/\lambda_i} - 1\right). \tag{15}$$

This expression allows us to compare V_1^{**} and V_2^{**} directly and determine which component to focus on, depending on the parameters (p_1, λ_1), (p_2, λ_2), and reward R. Although this expression technically solves the full problem, it is relatively unwieldy, as it includes endogenous V_j^* terms and fails to provide clean managerial insight on which component to pursue. Its limitation is also that in case there are n components rather than just 2, the algorithm complexity for finding the optimal component to focus on via this expression is $O(n \cdot 2^n)$,[1] making it computationally challenging in case of many components.

However, what becomes easy to derive from the expression is that in the context of the monkey-and-pedestal example, it is not universally true that one should focus on the monkey (the task with lower λ_i and p_i) first, which we formalize in the following corollary.

Corollary 1 *For every $\lambda_1, p_1, p_2, c, \delta, R$, there exists $\bar{\lambda}_2$ such that if $\lambda_2 > \bar{\lambda}_2$, then $V_2^{**} > V_1^{**}$.*

Thus, as long as there is at least a minuscule amount of uncertainty about whether one can successfully build the pedestal ($p_2 < 1$) and the pedestal can be built either fast enough ($\lambda_2 > \bar{\lambda}_2$) it is optimal to do the pedestal first. Intuitively, if there is some uncertainty about whether one can execute an easy task, it may be worth verifying that first, before embarking on the hard task.

[1] While, at first glance, evaluating this requires considering $n!$ options and then running a n-step recursion for each, resulting in complexity of $O(n \cdot n!)$, this can be improved to $O(n \cdot 2^n)$ by storing calculations that can be reused; the subset-based algorithm for doing so is given in Appendix EC.3.3

Note that there is a close relationship between the notions of speed and cost in this model, as λ_i also determines the expected costs to successfully complete a component. Thus, a high λ_i also implies a low expected cost.

From Corollary 1 it is clear that it is not necessarily optimal to work on the component with lower p_i; however, decreasing p_i does make a component more attractive to prioritize. We formalize this in the following Corollary.

Corollary 2 *If it is optimal to work on component* 1 *first, then it is still optimal to work on component* 1 *if* p_1 *is replaced by* $\tilde{p}_1 < p_1$.

While (15) does not provide a simple criterion for which component to pursue, we can derive some alternative representations which enable us to see which are the key tradeoffs here. Thus, we introduce the following change of variables:

$$x_1 = \frac{1}{p_1} - 1, \quad y_1 = \frac{1}{q_1} - 1, \quad x_2 = \frac{1}{p_2} - 1, \quad y_2 = \frac{1}{q_2} - 1,$$

where $q_i = 1 - c/\lambda_i R$. Similarly, define

$$y_1' = \frac{c}{\lambda_1 V_2^* - c}, \quad y_2' = \frac{c}{\lambda_2 V_1^* - c}.$$

Using these substitutions leads to the following compact expressions, where the optimal value functions are re-parametrized as functions of x_i, y_i and y_i':

$$\frac{V_i^*(x_i, y_i)}{R} = \frac{\frac{\lambda_i}{\lambda_i+\delta}\left(1 - (x_i y_i)^{1+\delta/\lambda_i}\right) + \frac{\lambda_i}{\delta} x_i y_i \left((x_i y_i)^{\delta/\lambda_i} - 1\right)}{(x_i+1)(y_i+1)}, \tag{16}$$

$$\frac{V_i^{**}(x_i, y_i')}{V_j^*} = \frac{\frac{\lambda_i}{\lambda_i+\delta}\left(1 - (x_i y_i')^{1+\delta/\lambda_i}\right) + \frac{\lambda_i}{\delta} x_i y_i' \left((x_i y_i')^{\delta/\lambda_i} - 1\right)}{(x_i+1)(y_i'+1)}. \tag{17}$$

Both V_i^* and V_i^{**} are strictly decreasing in both x_i and y_i (or y_i'), reflecting the diminishing expected payoff as either prior belief or project quality worsens. In the general case, we can see that the factors that influence which component to pursue first are: (a) Relative standalone payoffs: $V_1^* \sim V_2^*$, (b) Asymmetry in priors: $x_1 \sim x_2$, (c) Differences in project qualities: $y_1' \sim y_2'$ or equivalently, $\frac{V_2^*}{\lambda_2} \sim \frac{V_1^*}{\lambda_1}$, (d) The discounting effect through the factor $\lambda_i/(\lambda_i + \delta)$.

Due to the nonlinearity of these expressions and their countervailing effects, general analytical results are difficult to obtain. We will deal with this in two ways. First, in Section 4, we consider moonshot projects (ones where reward is very high) which are both of practical relevance and have clean analytical results. Then, in Section 5, we propose an approximate index policy for this problem which is intuitive and asymptotically optimal.

4. Moonshot Projects

The tech industry and startups often pursue projects that are unlikely to work out, but offer immense rewards compared to the costs in case of success; these are typically referred to as moonshot projects. We can consider the optimal component prioritization for that scenario by considering the limiting behavior of the model as R becomes very large.

Proposition 3 *There exists $\bar{R}$ such that if $R > \bar{R}$, the optimal component to work on is given by a lexicographic index policy: work on whichever component has a higher arrival rate (λ_i); if they are equal, work on the one with the lower prior probability of being feasible (p_i).*

This result implies that, in the context of moonshots, the optimal policy is to prioritize the component that can be completed more quickly *if* it turns out to be feasible—effectively disregarding (except as the last tie-breaking criterion) the actual likelihood that the component is feasible in the first place.

This conclusion is strikingly at odds with both the conventional wisdom to "work on the monkey before the pedestal" and the "fail fast, fail cheap" intuition underpinning Proposition 2. Under the conditions of Proposition 3, the optimal policy instead suggests beginning with the figurative pedestal—that is, the component with the higher λ_i. This is particularly noteworthy given that the original "monkey first" recommendation was conceived precisely with moonshot projects in mind. Of course, this result is contingent on (a) there being at least some uncertainty about feasibility of the pedestal ($p_i < 1$), and (b) the pedestal being easier to build than the monkey is to train (higher λ_i). If the pedestal is a project with no uncertainty at all, or one that is significantly more expensive than the monkey, the original prescription to do the monkey first holds. To better understand this divergence, we now examine the underlying intuition—both mathematical and managerial.

From a mathematical perspective, consider the asymptotic behavior of the value function under different prioritization policies. As established in equation (15), we obtain

$$\lim_{R\to\infty} \frac{dV_1^{**}}{dR} = \lim_{R\to\infty} \frac{dV_2^{**}}{dR}$$

That is, both prioritization strategies yield the same asymptotic slope. This outcome is intuitive: as R grows large, the marginal cost becomes negligible, the likelihood of abandoning the project vanishes, and the expected reward converges to the ex-ante probability that both components are feasible—namely, $p_1 p_2$ times the expected discount factor for when the reward is obtained.

What distinguishes the policies in this limit, then, is not where the value functions converge, but *how quickly* they converge. The rate of convergence is driven by how rapidly the uncertainty about feasibility is resolved, which in turn is governed by the arrival rate λ_i. From a managerial standpoint, the reason λ_i dominates in this regime is that the long-run value of the project depends solely on the joint feasibility probability $p_1 p_2$, while λ_i determines how quickly the firm learns about feasibility. Thus, faster information acquisition—enabled by a higher arrival rate—becomes the key determinant of optimal action under high-reward or low-cost scenarios.

Figure 1 illustrates the exact two-component comparison. The decision-maker chooses among three possibilities: abandon immediately, prioritise component 1, or prioritise component 2. The unlabelled region corresponds to immediate abandonment, while the active region is divided by the indifference curve $V_1^{**} = V_2^{**}$.

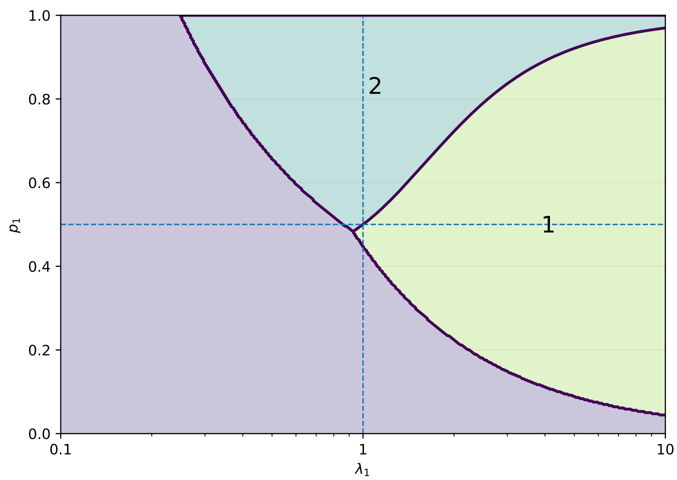

Figure 1 Priority and abandonment regions. The figure fixes $p_2 = 0.5$, $\lambda_2 = 1$, $R = 10$, $c = 1$, and $\delta = 0.5$, and varies p_1 and λ_1. The unlabelled region corresponds to immediate abandonment, where $\max\{V_1^{**}, V_2^{**}\} \leq 0$. In the active region, the boundary between regions labelled 1 and 2 is the indifference curve $V_1^{**} = V_2^{**}$. Regions labelled 1 and 2 indicate that it is optimal to prioritise component 1 and component 2, respectively. The figure shows that priority depends jointly on feasibility and speed of learning: a component with lower feasibility probability may be attractive to test first, but sufficiently fast experimentation can reverse the simple monkey-first logic.

This exact comparison becomes less convenient as the number of components grows. We therefore next develop a component-by-component index rule.

5. Approximate Index Policy

5.1. The Gittins Variant Index

Here, we propose an index policy that is easy to compute because each component's index depends only on that component's own primitives. The idea behind the policy is simple and consistent with the well-known Gittins index (Gittins 1979): Can we choose the optimal component based on only the characteristics of that component? More precisely, what would be the value function if we prioritize component i and work on it for a time T_i^*, that is exactly the same time as would be optimal to work on it if that was the only component remaining? Can we just prioritize the component for which this value function is the highest?

Much like in the case of applying the Gittins index to the Pandora's problem of Weitzman (1979) with unknown distributions, prioritizing components like this leads to the truly optimal policy in some simpler versions of the model,[2] and while it is not optimal in the general case, we will show that it still provides a good approximation (Beyhaghi and Cai 2024).

The value functions under this policy are $V(T_1^*,0)$ and $V(0,T_2^*)$, as given by (13). With some algebraic manipulation, the comparison between these two value function reduces to a relatively simple index given by the following proposition.

Proposition 4 (Gittins Variant Index) *Denote by $\hat{c}_i$ the expected cost of a decision maker who works on component i for time T_i^* or until the component is successfully completed, whichever comes first. Similarly, denote by $\hat{r}_i$ the expected discounted probability that such a decision maker successfully completes component i. Then, $V(T_1^*,0) \geq V(0,T_2^*)$ if and only if $I_1 \geq I_2$, where*

$$I_i = \frac{1-\hat{r}_i}{\hat{c}_i}.$$

In terms of model primitives, this index can be expressed as

$$I_i = \frac{\lambda_i + \delta - \lambda_i p_i \left(1 - e^{-(\lambda_i+\delta)T_i^*}\right)}{c\left[p_i\left(1-e^{-(\lambda_i+\delta)T_i^*}\right) + (\lambda_i+\delta)(1-p_i)\frac{1-e^{-\delta T_i^*}}{\delta}\right]}. \tag{18}$$

The index is relatively intuitive: putting aside the full formality of definitions in the proposition, one can think about it as failure probability divided by expected costs. The index also preserves the approach of "fail fast, fail often, fail cheap" that underlies Proposition 2. Faster arrival rate and lower probability of success both make the project more attractive to focus on (by decreasing $\hat{r}$), while higher expected cost makes it less attractive to focus on. In other words, the index is focused on ensuring that if the (hidden) state of the world is such that the project is inevitably going to be terminated, that failure state will be identified as quickly and cheaply as possible.

However, such an approach is not fully optimal, so we proceed to explore how good of an approximation it really is in the following proposition.

Proposition 5 (Performance Gap) *Denote by $V(I)$ the value function of a decision maker who follows the index policy of Proposition 4: they select the component with the highest index I_i to*

[2] If we have a discrete time-model, where one can only work on a component a single time, the Gittins variant index is truly optimal. Other than the discrete time, what makes this variant simpler is that lack of success is completely informative about feasibility of the component.

work on, and stop at time T_i^* in the absence of a success. Then, denoting by i the component with the higher index, we have

$$V_i^{**} - V(I) = c\left[p_i \frac{\delta e^{-(\delta+\lambda_i)T_i^{**}} + \lambda_i e^{-T_i^*(\delta+\lambda_i)} - (\delta+\lambda_i)e^{-\delta T_i^{**} - \lambda_i T_i^*}}{\delta(\delta+\lambda_i)} + (1 - p_i + p_i e^{-\lambda_i T_i^*})\frac{e^{-\delta T_i^{**}} - e^{-\delta T_i^*}}{\delta}\right] - V_j^*\left[p_i \frac{\lambda_i\left(e^{-T_i^{**}(\delta+\lambda_i)} - e^{-T_i^*(\delta+\lambda_i)}\right)}{\delta+\lambda_i}\right]. \tag{19}$$

Both terms in square brackets above are positive.

Furthermore, denoting by $B(I)$ the regret (the ex-post payoff loss from following I_i instead of the optimal policy on any sample path), it is bounded by

$$B(I) \leq c\frac{e^{-\delta \min\{T_i^{**}, T_j^{**}\}} - e^{-\delta(T_i^* + T_j^*)}}{\delta}. \tag{20}$$

The interesting property here—which should be evident from the absence of R in (20)—is that in any situation where the optimal policy would successfully complete the project, the index policy will do so as well. So, the index policy does not miss out on any potential upside of this project. However, it can have inefficiently high costs due to working longer than optimal on the component it prioritizes. Perhaps unintuitively, there are sample paths where misordering works out for the better, so it is not sufficient to just consider the case with misordering in derivation of (20). Yet, both of these issues of the index policy vanish when the reward for completing the project is sufficiently high, which we formalize in the following Proposition.

Proposition 6 (Asymptotic Optimality) *As rewards increase, this index policy is asymptotically optimal and the regret of using the index policy vanishes, i.e.,* $\lim_{R\to\infty} V(I) = \lim_{R\to\infty} \max_i\{V_i^{**}\}$ *and* $\lim_{R\to\infty} B(I) = 0$*. Furthermore, there exists $\bar{R}$ such that for all $R > \bar{R}$, the index policy will select the correct component to focus on (the same one as the fully optimal policy).*

The property that value functions converge as $R \to \infty$ is not surprising: essentially any policy that does not give up on working on the project will have an unbounded expectation when reward gets infinite as any costs incurred become negligible compared to the reward for successful completion.

That regret bound vanishes is somewhat more insightful, as it also means that as the reward gets large, the index policy will also not experience inefficiencies on any sample path, not just in expectation.

Arguably, the most interesting part of the proposition though, is that for sufficiently high reward levels (and not just in the limit) the index policy will correctly prioritize the components. This motivates a slightly modified version of the index policy that is fully optimal for sufficiently high values of R, as we formalize in the following result.

Corollary 3 *Denote by $\tilde{V}(I)$ the value function of a decision maker who first prioritizes the component i with the highest index I_i, and then follows the corresponding optimal stopping policy when component i is chosen first, yielding value V_i^{**} as given by (15). If $R > \bar{R}$, where $\bar{R}$ is the threshold from Proposition 6, then this decision maker behaves optimally, in the sense that*

$$\tilde{V}(I) = \max_k V_k^{**}.$$

While the modified index policy is fully optimal for high reward values, and has higher expectation than the original index for any parameter values, it does lose some beneficial properties of the original. Namely, as the modified index policy does not necessarily explore both components for longer than the optimal policy, it can miss the reward in some cases where the optimal policy would have found it. As a result, its regret can actually be higher than the original index.

Noting that (19) established in Proposition 5 is also an upper bound for the error of the modified index policy, we also provide here an illustrative numerical example of how close to the true optimum can the modified index policy reasonably be expected to be.

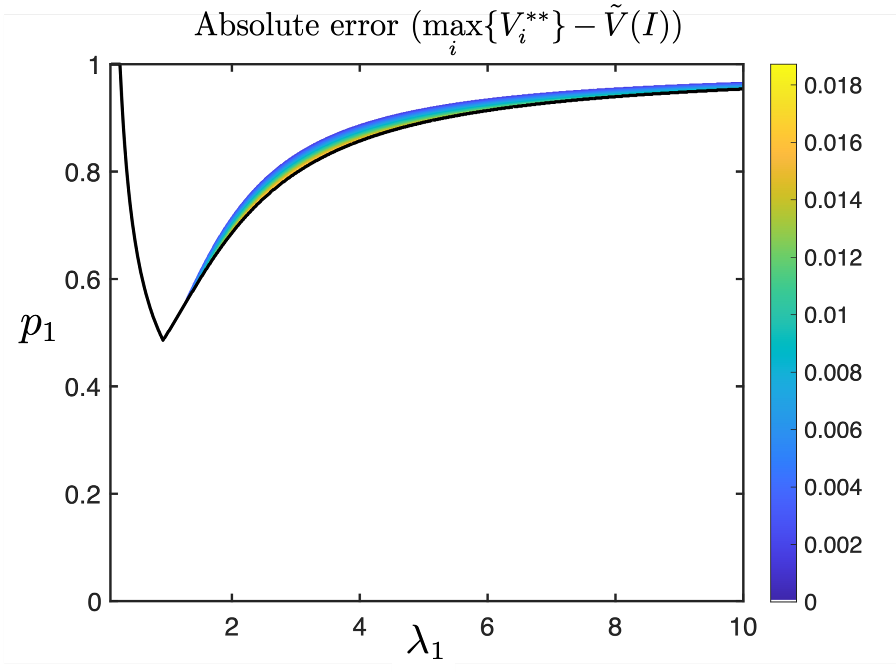

Figure 2 **The absolute error under the modified index policy, under parameters $R = 10$, $p_2 = 0.5$, $\lambda_2 = 1$, $c = 1$, $\delta = 0.5$. The black line is the indifference curve where the optimal policy is indifferent between which component to focus on.**

Here, it can be observed that for vast majority of parameter values (the entire white region of the figure), the modified index policy will be fully optimal. This happens whenever the index selects the correct component to focus on, as the modified index policy makes no timing errors. The lack of timing errors is also the reason why the modified index policy makes no errors for very low values of λ_1, as under those values it correctly identifies that it is optimal to abandon the project right away.

All errors occur close to the indifference curve between the two components (the region with errors is the colored part of the figure), as that is the region where the index will make ordering errors. Yet, this proximity to indifference curve also causes the errors to be fairly small as it is also the region where the value function under different orderings are relatively close, with the difference between them vanishing at exactly the indifference curve (the largest possible error in the example here is less than 20% of a single-period cost of working on a component). Thus, it may be reasonable to expect that the index will exhibit good performance even in the small-R case. The next subsection is devoted to testing this index both analytically and numerically.

5.2. Performance of the index for small rewards

In this subsection we show that a sharp *finite-R* bound on the error of the modified index policy is available, and that it implies a clean parameter-explicit *sufficient condition* for exact optimality that does not require R to be large. Consistently with the example in Figure 2, the picture that emerges is that the modified index policy fails to be exactly optimal only on a thin band around the indifference curve between the two components, and that the size of the failure on this band is controlled by a function of single-component primitives that we make explicit below.

We first derive an algebraic decomposition of the value gap. Throughout this subsection we restrict attention, without loss of generality, to the active region $p_i R\lambda_i > c$.[3] Define the following three objects:

$$\Delta := V_1^{**} - V_2^{**}, \tag{21}$$

$$\Phi := V(T_1^*, 0) - V(0, T_2^*), \tag{22}$$

$$O_i := V_i^{**} - V(T_i^*, 0) \geq 0, \quad i \in \{1, 2\}. \tag{23}$$

The interpretation is direct: Δ is the gap on which the *optimal* policy bases its sequencing decision (it works on component 1 first only if $\Delta \geq 0$); Φ is the gap on which the *index* bases its decision, since by Proposition 4 the equivalence $V(T_1^*, 0) \geq V(0, T_2^*) \Longleftrightarrow I_1 \geq I_2$ holds, so $\text{sign}(\Phi) = \text{sign}(I_1 - I_2)$; and O_i is the loss the decision-maker incurs from continuing to work on component i for time T_i^* rather than the truly optimal time T_i^{**}, given that the relevant continuation reward is $V_j^* < R$.

[3] If this is not satisfied, it is optimal to immediately stop, consistent with the modified index policy, leading trivially to an uninteresting situation with no optimality gap.

Non-negativity of O_i follows because T_i^{**} is the maximizer of the single-component value function with reward V_j^*. Moreover, the right-hand side of Eq. (19), with i replaced by either 1 or 2, is exactly O_i. The starting point of our analysis is the identity below, which follows immediately from $V_i^{**} = V(T_i^*, 0) + O_i$:

$$\Delta \;=\; \Phi \;+\; (O_1 - O_2). \tag{24}$$

The decomposition splits the optimal sequencing gap into two pieces: the index's own sequencing gap Φ, and an asymmetry term $O_1 - O_2$ that measures how much more inefficient overshooting is on component 1 relative to overshooting on component 2.

The modified index policy of Corollary 3 prioritizes the component i with the higher I_i and then plays optimally on it, attaining V_i^{**}. Its absolute error is therefore

$$\max_k V_k^{**} - \tilde{V}(I) \;=\; \begin{cases} 0 & \text{if } \operatorname{sign}(\Phi) = \operatorname{sign}(\Delta), \\ |\Delta| & \text{otherwise.} \end{cases}$$

Combining this with (24) gives the following finite-R bound. For all parameter values in the active region, the absolute error of the modified index policy of Corollary 3 satisfies

$$\max_k V_k^{**} \;-\; \tilde{V}(I) \;\leq\; \max\bigl\{0,\; |O_1 - O_2| - |\Phi|\bigr\}, \tag{25}$$

with equality whenever the index makes an ordering error, i.e., whenever $\operatorname{sign}(\Phi) \neq \operatorname{sign}(\Delta)$. In particular,

$$|\Phi| \;\geq\; |O_1 - O_2| \quad\Longrightarrow\quad \tilde{V}(I) = \max_k V_k^{**}, \tag{26}$$

i.e., whenever the index gap $|\Phi|$ exceeds the asymmetry $|O_1 - O_2|$ of the overshoot losses, the modified index policy is fully optimal.

The geometry of the errors warrants a brief discussion. The set of parameter points on which $\operatorname{sign}(\Phi) \neq \operatorname{sign}(\Delta)$ is the symmetric difference of the optimal indifference curve $\{\Delta = 0\}$ and the index indifference curve $\{\Phi = 0\}$. Both are codimension-one continuous surfaces in $(p_1, \lambda_1, p_2, \lambda_2, c, \delta, R)$-space, and they coincide on the locus $\{O_1 = O_2\}$. Equations (25)–(26) formalize the visual observation in Figure 2: errors sit in a thin band that pinches the two indifference curves together on the locus where the overshoot asymmetry vanishes.

We now turn our attention to numerical results on the performance of the modified index for small values of R. From equation (25), the performance gap exists only when $|O_1 - O_2|$ is greater than $|\Phi|$, which is actually a rarity across the parameter space.

We check this numerically across a wide parameter set. For each value of R in a logarithmic grid spanning over four orders of magnitude, $R \in \{0.1, 0.2, 0.5, 1, 2, 5, 10, 25, 50, 100, 500, 5{,}000\}$, we sampled $n = 30{,}000$ parameter points $(p_1, p_2, \lambda_1, \lambda_2, c, \delta)$ as $p_i \sim \mathcal{U}(0.01, 0.99)$, $\log \lambda_i \sim$

$\mathcal{U}(\log 0.05, \log 20)$, $\log c \sim \mathcal{U}(\log 0.05, \log 20)$, $\delta \sim \mathcal{U}(0.02, 3.0)$. The log-uniform draws on λ_i and c ensure that ratios of arrival rates and the cost-to-reward ratio are well-behaved over several orders of magnitude. For each point we restrict to the active region $\max_k V_k^{**} > 0$ and record (i) the relative error $(\max_k V_k^{**} - \tilde{V}(I))/\max_k V_k^{**}$, (ii) whether the sufficient condition $|\Phi| \geq |O_1 - O_2|$ of (26) holds, and (iii) whether the modified index is exactly optimal.

Table 1 summarizes the Monte Carlo experiment. The picture is uniform across rewards and qualitatively unchanged by the broadening of the parameter ranges.

		Relative error					
R	#active	mean	90th %	99th %	99.9th %	%exact	%suff
0.1	418	0.07%	0.0%	0.0%	15.6%	99.5	98.1
0.2	897	0.03%	0.0%	0.0%	8.5%	99.2	97.4
0.5	2,062	0.13%	0.0%	0.0%	27.4%	99.0	94.8
1	3,333	0.14%	0.0%	0.1%	37.5%	98.9	92.7
2	5,209	0.13%	0.0%	0.1%	41.4%	98.7	88.6
5	8,381	0.14%	0.0%	0.0%	61.3%	98.9	82.7
10	11,080	0.11%	0.0%	0.0%	48.2%	98.7	77.5
25	14,747	0.09%	0.0%	0.0%	25.8%	98.6	73.3
50	17,839	0.08%	0.0%	0.0%	20.6%	98.4	69.5
100	20,450	0.09%	0.0%	0.0%	31.2%	98.3	66.8
500	25,631	0.08%	0.0%	0.0%	16.2%	97.6	62.0
5,000	29,005	0.04%	0.0%	0.0%	6.5%	97.5	59.5

Table 1 Relative error of the Monte Carlo simulation of the modified index policy under $30{,}000$ samples per R. "#active" counts the parameter points on which $\max_k V_k^{**} > 0$. "%exact" is the fraction of active points on which the modified index attains the optimum; "%suff" is the fraction on which the sufficient condition $|\Phi| \geq |O_1 - O_2|$ holds.

Three features stand out. First, *mean and high-quantile performance is essentially exact.* At every R tested, the mean relative error stays below 0.15% and the 90th and 99th percentiles round to zero. The fraction of active points on which the modified index is exactly optimal never drops below 97.5%.

Second, the *sufficient condition* $|\Phi| \geq |O_1 - O_2|$ *decays smoothly with* R but remains the main guarantee even at large R: it holds on 98% of active points at $R = 0.1$, and on 60% at $R = 5{,}000$. The gap between "%exact" and "%suff" measures how often the policy is exactly optimal *despite* failing the sufficient condition—i.e., how often the sign of Φ agrees with the sign of Δ even though $|\Phi|$ is the smaller of the two terms in (24). This gap grows with R, consistent with the asymptotic mechanism of Proposition 6: as R grows the overshoots O_1 and O_2 scale similarly, so $|O_1 - O_2|$ stays modest while neither O_i shrinks individually, and the index's sign continues to align with Δ.

Third, the apparently alarming far-tail entries—99.9th percentiles running up to 61%, with isolated maxima of 100% in the full active set—are diagnostic of a single, well-understood failure mode and not of any breakdown of the policy. Inspection of the worst-case scenarios returned by the Monte Carlo simulation confirms that every such case has V_{opt} on the order of 10^{-4} to 10^{-3}, i.e., the project is at the very edge of the active region. On these points the absolute error is at most 10^{-3} but the relative error can be arbitrary because the optimal value is itself near zero. To control for this we recomputed the same statistics on the “robust” subset $V_{\mathrm{opt}} > 0.05\,c$ (See Figure 3, or for more details, Table EC.1 in the Appendix). The 99.9th-percentile relative error on the robust subset stays below 11% for $R \le 50$ and below 7% for $R \le 2$, with no substantively large values anywhere.

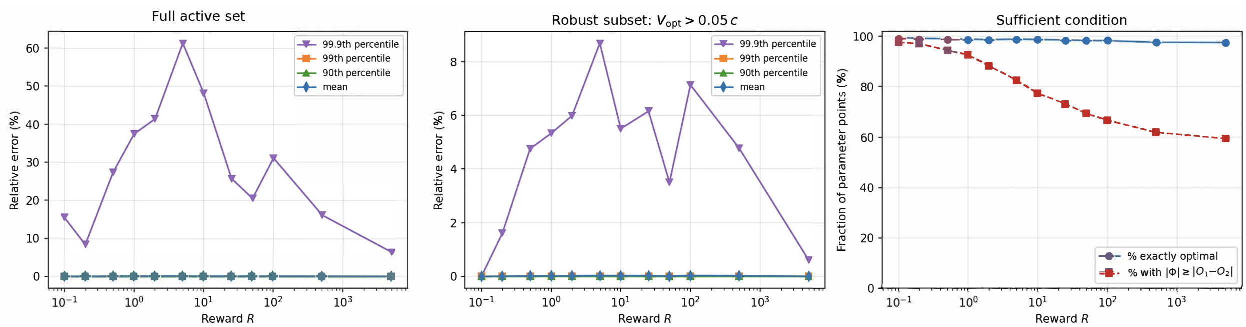

Figure 3 **Modified-index policy under** $30{,}000$ **parameter points per value of** R**. Left: relative error on the full active set** $\max_k V_k^{**} > 0$**. The large 99.9th-percentile and maximum values at small** R **are driven by parameter points near the boundary of the active region, where** V_{opt} **is itself near zero. Middle: same statistics on the robust subset** $V_{\mathrm{opt}} > 0.05\,c$**. Right: fraction of active points on which the modified index is exactly optimal (blue, solid), and the fraction on which the sufficient condition** $|\Phi| \ge |O_1 - O_2|$ **holds (red, dashed).**

In summary, combined with the closed-form bound of (25) and the moonshot result of Proposition 6, this gives the modified index a strong end-to-end story. Asymptotically, the policy is exactly optimal. At any finite R, it is exactly optimal on at least 97.5% of the parameter space; it is provably exactly optimal on a much larger and parameter-explicit $|\Phi| \ge |O_1 - O_2|$ set; and where it is not optimal, its absolute error is uniformly bounded by $|O_1 - O_2| - |\Phi|$, an object that the practitioner can compute in closed form from single-component primitives.

6. General Information Structure

Much of the paper considers a specific information structure for the arrival rate of success when working on a component i: either the component is infeasible (in which case $\Lambda_i = 0$) or the component is feasible with a fixed, known arrival rate ($\Lambda_i = \lambda_i > 0$). Here we return to the most general version of the model, and allow for any arbitrary prior for the arrival rate. However, this will result in more complexity and different results. Most notably, the condition of Theorem 2 does not universally hold for all information structures (thus non-corner solutions are possible), which we illustrate with the following example.

Consider the case of a simple two-point probability distribution where the uncertainty is *not* about feasibility, i.e., $\Lambda_i \in \{\lambda_{iA}, \lambda_{iB}\}$, where $\lambda_{iA}, \lambda_{iB} > 0$. Thus, there is no question of whether it is possible to complete the project, just how fast (and thus at which cost) can it be completed.

In this situation, it can easily become optimal for the decision-maker to embark on completing the project if they are sufficiently confident at least one of the components has a fast arrival rate, in which case the optimal policy has an interior solution: work on one component until either a success is encountered or (due to lack of success) the posterior probability of that component having a fast arrival rate becomes sufficiently low, at which point switch to exploring the other component. Figure 4 illustrates one such case, with interior solution $T_1 \approx T_2 \approx 0.15$.

Figure 4 **Topographical plot of** $V(T_1, T_2)$ **for parameter values** $p_1 = p_2 = 0.3, R = 2.5, c = 1$**,** $\delta = 0$**,** $\lambda_{1A} = \lambda_{2A} = 0.5$ **and** $\lambda_{1B} = \lambda_{2B} = 10$**.**

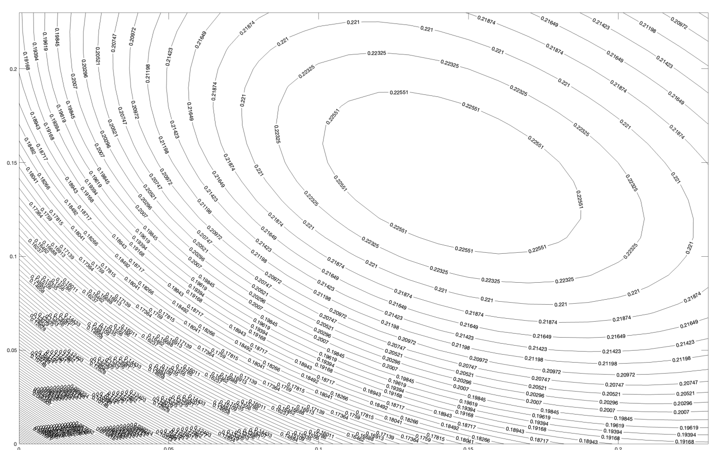

This example underscores the importance of the nature of uncertainty. When uncertainty pertains to whether a component is feasible, as in our baseline model, the optimal policy is invariably a corner solution: the decision maker devotes attention to a single component until either a success is realized or lack of arrivals makes the posterior sufficiently pessimistic that termination of the entire project is optimal. This form of uncertainty is well suited to projects involving novel and highly complex technologies—such as the Hyperloop, fusion energy, the James Webb Space Telescope, or Neuralink—where it is unclear ex ante whether successful development is feasible.

A fundamentally different setting arises when the decision maker knows that successful completion is possible, but faces uncertainty about the time or cost required to achieve it. The example above falls into this second category. In this case, the model parameters may imply that pursuing the project is worthwhile only if at least one component can be developed sufficiently quickly and at sufficiently low cost. As a result, an interior solution may become optimal: the decision maker initially works on one component, but if no success is observed soon enough—thereby generating a sufficiently unfavorable posterior about its arrival rate—switches to the other component to explore whether that component can be completed more efficiently. This scenario fits well with the companies that are adopting emerging technologies for which feasibility has already been established, but there is still uncertainty about whether they can implement those technologies at a sufficiently low cost to be competitive in the market. For example, Mercedes Benz moving production to electric vehicles, or Anthropic entering the market for fundamental large language models may fit this situation well.

More broadly, our analysis shows that the optimal prioritization of components in multi-component projects can differ sharply depending on whether the underlying uncertainty concerns feasibility or execution cost.

From the methodological standpoint, presence of interior solution is challenging, since a lot of our results rely on Theorem 2. We approach this challenge as follows: First, in Proposition 7 we show that the conditions of Theorem 2 are universally satisfied for moonshot projects, i.e., ones where the reward is sufficiently high and there is at least some uncertainty about feasibility. Second, in Proposition 9 we show that the index policy is still asymptotically optimal even in the case with a general prior, without the need for conditions of Theorem 2. We proceed with the statement of the first of these results.

Proposition 7 (Corner solutions, general prior with feasibility uncertainty) *Suppose for $i = 1,2$ that $\mathbb{P}(\Lambda_i = 0) > 0$; $\lambda_{i1} := \inf\big(\mathrm{supp}(\Lambda_i) \cap (0,\infty)\big) > 0$; λ_{i1} is an isolated point of $\mathrm{supp}(\Lambda_i) \cap (0,\infty)$, i.e., either $\mathbb{P}(\Lambda_i = \lambda_{i1}) > 0$ or $\inf\big(\mathrm{supp}(\Lambda_i) \cap (\lambda_{i1},\infty)\big) > \lambda_{i1}$; $\mathbb{E}[\Lambda_i] < \infty$.*

Then there exists $\bar{R}$ such that for all $R > \bar{R}$, every optimum of (14) is a corner solution.

The proposition extends the corner-solution conclusion of Proposition 2, which was confined to two-point priors, to all moonshots under the full generality of priors that combine a positive atom at zero with any positive support with an isolated low point λ_{i1}. In economic terms, the first structural assumption is that there is at least *some* prior chance the component is infeasible ($\mathbb{P}(\Lambda_i = 0) > 0$); the lower bound $\lambda_{i1} = \inf(\mathrm{supp}(\Lambda_i) \cap (0,\infty)) > 0$ is a mild regularity condition that excludes the technically degenerate case in which arbitrarily small positive arrival rates are possible.

The assumption that λ_{i1} is an isolated point is a technical assumption needed to guarantee the asymptotic concentration of the positive-part posterior on λ_{i1} to be at the exponential rate required by the proof. Two salient cases satisfy the assumption automatically: (i) Two-point priors (Proposition 2) and any finite discrete prior have an atom at λ_{i1}. (ii) Priors of the form "a positive mass at some specific positive rate plus an arbitrary continuous tail above" (the canonical setting for feasibility-with-known-completion-rate models) again have an atom at λ_{i1}.

The case ruled out by the isolated point assumption is a prior whose positive support is a single interval $[\lambda_{i1}, \bar{\lambda}]$ with continuous and positive density at λ_{i1}^{+}. In this case the conclusion of the proposition can fail: a calculation (see Appendix EC.3.4) shows that the posterior variance on the positive part decays only polynomially in t_i (specifically as $1/t_i^2$), the slope of $V_i^*(t_i)/\hat{\Lambda}_i(t_i)$ becomes positive at the interior first-order condition for R large, and the Hessian determinant becomes positive—in other words, an interior local maximum can exist.

Under these assumptions, the moonshot regime (R large) makes interior solutions impossible: for every R above some threshold $\bar{R}$, the decision-maker focuses on a single component until either a success arrives or the posterior on that component is sufficiently pessimistic to abandon the project entirely, switching to the other component only upon success. This is the same "fail fast, fail cheap" qualitative prescription that drives Proposition 2, now robust to whatever finer structure the positive part of the prior on Λ_i may have. In particular, the result confirms that the two-point analysis of the earlier sections is not an artifact of the two-point reduction: the corner-solution conclusion is the rule, not the exception, whenever the moonshot reward dominates and feasibility is genuinely uncertain.

Thus, it may be reasonable to expect that the approximate index policy of Section 5 will also have good performance in the setting and asymptotic optimality even under general distributions with feasibility uncertainty. We start with the following proposition that identifies the behavior of the index policy for cases of high reward (moonshots).

Proposition 8 (General index for moonshots) *Assume for $i=1,2$ the distribution of Λ_i satisfies the assumptions of Proposition 7. There exists $\bar{R}$ such that for all $R>\bar{R}$, the index of Proposition 4 obeys:*

1. *if $\lambda_{i1}\neq\lambda_{j1}$, the index selects the component with the* higher *of $\lambda_{i1},\lambda_{j1}$;*
2. *if in addition the distribution of each Λ_i places a positive atom $\pi_{i1}:=\mathbb{P}(\Lambda_i=\lambda_{i1})>0$ at λ_{i1} and $\lambda_1:=\lambda_{i1}=\lambda_{j1}$, the index selects the component with the higher value of*

$$\frac{\pi_{i0}}{\mathbb{E}_{\Lambda_i}\left[\frac{1}{\Lambda_i+\delta}\right]}\left(\frac{c\,\pi_{i0}}{\pi_{i1}\,\lambda_1}\right)^{\delta/\lambda_1}. \tag{27}$$

Part 1 says the moonshot index is essentially *lexicographic* in λ_{*1}: it always favors the component whose smallest plausible arrival rate is higher—the one whose feasibility (or infeasibility) is resolved faster in the absence of arrivals. Part 2 records that when $\lambda_{i1} = \lambda_{j1}$, unlike Proposition 3, the policy will no longer be a lexicographic one, but will depend on a non-trivial interaction of several model parameters, as seen in (27). The choice is more nuanced, but the discrete tie-breaker from the original version carries in essence over to the setting when both priors place mass at the common λ_1; the expression (27) is large when π_{i0} is large (high prior chance of infeasibility) and when π_{i1} is small, both of which speed up the resolution of feasibility uncertainty.

An interesting property is that for high rewards, this causes the policy to depend on the underlying distribution only via its expectation, and the probability masses on the two smallest elements that have a mass (or the infimum of the non-zero part of the support, in case only 0 has mass). All other elements of the support matter only through their effect on the expectation. There is reasonable intuition for this behavior: when rewards are high, it will be optimal to work on a component relatively long, even in the absence of success. However, after a long time working a component with no success, the posterior distribution for Λ_i will increasingly resemble a two-point distribution with support $\{0, \lambda_{i1}\}$. This result enables us to establish the full asymptotic optimality of the index policies for general distributions:

Proposition 9 (Asymptotic optimality of the index policy under feasibility uncertainty) *Suppose that for $i = 1, 2$, the distribution of Λ_i satisfies the assumptions of Proposition 7. Denote by V^{**} the value function under the optimal policy and by $V(I)$ and $\tilde{V}(I)$ the value functions under the index policy of Proposition 4 and the modified index policy of Corollary 3, respectively. Then:*

1. *Both the index policy and the modified index policy are asymptotically optimal:*
$$\lim_{R\to\infty} \frac{V(I)}{V^{**}} = 1 \qquad \text{and} \qquad \lim_{R\to\infty} \frac{\tilde{V}(I)}{V^{**}} = 1.$$
2. *If $\lambda_{i1} \neq \lambda_{j1}$, there exists $\bar{R}$ such that for all $R > \bar{R}$, the index policy selects the same component as the optimal policy. Consequently, the modified index policy is fully optimal: $\tilde{V}(I) = V^{**}$ for all $R > \bar{R}$.*

Part 1 records the universal statement: under any prior with positive mass on infeasibility and a positive gap above zero in the support, both the index policy and the modified index policy attain the same asymptotic value as the optimal policy. In particular, the modified index policy remains asymptotically optimal even on the boundary case $\lambda_{i1} = \lambda_{j1}$, where which component the index actively picks may not coincide with the optimum; the leading-order value is the same on both sides, so picking the wrong side incurs only a vanishing relative loss. Part 2 strengthens this from a limit to a strict claim at finite (though large) R: when the smallest plausible arrival rates of the two components differ, the index policy not only matches the optimal asymptote but actively picks the same component.

7. Generalization to more than two components

For expositional clarity the body of the paper has focused on the $n = 2$ case. In this section we record which of the paper's results extend to a project consisting of an arbitrary number $n \geq 2$ of independent components. The picture is clean: the entire corner-solution backbone of the paper carries over under the same slope condition, and the only results that remain two-component-specific are a small collection of closed-form objects whose explicit formula depends on $n = 2$.

Throughout this section, let $\mathbf{S} \subseteq \{1, \ldots, n\}$ index the components not yet attempted, and let $V^*(\mathbf{S})$ denote the optimal value of the remaining sub-project once the outcomes on $\{1, \ldots, n\} \setminus \mathbf{S}$ have been resolved. By the same recursive logic as in the two-component case, V^* satisfies the dynamic programming recursion

$$V^*(\mathbf{S}) = \max\Big\{0,\ \max_{i \in \mathbf{S}} V_i^*\big(V^*(\mathbf{S} \setminus \{i\})\big)\Big\}, \qquad V^*(\emptyset) = R, \tag{28}$$

where $V_i^*(\rho)$ denotes the single-component optimal value of working on component i with continuation reward ρ. Computing $V^*(\{1, \ldots, n\})$ from (28) bottom-up over subsets, with the closed-form V_i^* at each step, requires $O(n \cdot 2^n)$ work—a significant improvement over the $O(n \cdot n!)$ cost of enumerating all permutations (see Appendix EC.3.3 for the exact $O(n \cdot 2^n)$ algorithm). The natural extension of the modified index policy of Corollary 3 is to compute the index of Proposition 4 for each component (only once, under the prior, so the ranking is not updated after intermediate successes), after which the order in which the DM works on the components is in the decreasing order of their index; this is computed in $O(n)$ time.

Most of the paper's results extend to this setting. The model definition, the control-space collapse of Theorem 1, and the single-component analysis of Proposition 1 involve no two-component-specific structure and carry over verbatim, with the only change being that the state (T_1, T_2) is replaced by an n-tuple $(T_1, \ldots, T_n)$. The construction of the index I_i in Proposition 4 also depends only on single-component primitives and is unchanged.

The corner-solution backbone of the paper extends under the same slope condition: the n-component analogue of Theorem 2 holds under the (strict) slope condition $V_i^*(t_i)/\hat{\Lambda}_i(t_i)$ strictly decreasing in t_i for every component. While this does require an expanded proof, we prove this explicitly in Theorem 3 (in Appendix EC.3.1) by a direct Hessian analysis at an interior critical point, identifying a vector $v = e_i - e_{\sigma(i)} \in \mathbb{R}^n$ with $v^\top H_n v > 0$ via the inductive corner-preference function σ together with a pigeonhole on the dispersion ratios $\mathrm{Var}(\Lambda_k)/\hat{\Lambda}_k$. As immediate corollaries, the corner-solution results of Proposition 2 (two-point priors) and Proposition 7 (general priors with feasibility uncertainty at large R) extend verbatim to arbitrary n: the slope condition needed at each component in both proofs is per-component and is preserved under any number of remaining components.

The moonshot characterization of Proposition 3 similarly extends, with the same lexicographic structure: an expansion of V_i^{**} around $R=\infty$ shows that the optimal first component is the one with the highest λ_{i1}, and the tie-breakers cascade recursively to subsequent components by the same argument used for $n=2$. The asymptotic-optimality results of Proposition 6 and its general-prior counterpart Proposition 9 extend in both parts: the leading R-coefficient of $V^*(\{1,\dots,n\})$ equals $R\prod_k \hat{r}_k^\infty$ regardless of the order in which components are sequenced (so $V(I)/V^* \to 1$ and $\tilde{V}(I)/V^* \to 1$ unconditionally), and the "picks-the-correct-component" clauses follow from the n-component no-switching just established together with the lexicographic structure of the index, which is itself a per-component asymptotic. Finally, the full-optimality conclusion of Corollary 3 and Part 2 of Proposition 9 extend in the same way.

The only results that do not extend are closed-form expressions in two-component-specific form. Independent of the analytical extension just described, a small set of results in the paper are closed-form objects defined on pairs of components and have no single-formula analogue for $n>2$. The performance gap of (25) is one such: its expression compares V_i^{**} to the value of the index policy via two pairwise terms with no obvious n-component reduction. The closed-form finite-R bound of Proposition 5, $\max_k V_k^{**} - \tilde{V}(I) \le \max\{0, |O_1 - O_2| - |\Phi|\}$, is likewise specific to the two-component decomposition $\Delta = \Phi + (O_1 - O_2)$: for $n > 2$ the analogous decomposition replaces Φ and the single pair (O_1, O_2) by an n-component ranking object and the resulting expression is no longer a single scalar comparison. These results are best understood as sharp closed-form refinements of the general theory that happen to take their cleanest form at $n=2$.

Lastly, there is strong numerical evidence for good performance of the index policy even in a small-R setting for $n>2$. While the analytical extensions described above settle the corner-solution and asymptotic-optimality questions at arbitrary n, the practical performance of the modified index policy is corroborated by Monte Carlo simulation. Appendix EC.3.2 reports a study comparing the policy to the exact $O(n\cdot 2^n)$ subset-DP optimum on random instances with n ranging from 2 to 12. Mean relative error grows smoothly with n at fixed R, mean absolute error saturates at minute amounts (order of magnitude 10^{-3}); the reason for this arguably spectacular performance in terms of absolute errors is that the vast majority of index errors occur when the project is near the boundary of not being even worth pursuing, and in such cases, even if the relative error is large, the absolute error is minuscule. The moonshot-regime asymptotic optimality survives intact: at R large the policy is essentially exact for every n we tested. Combined with the analytical extensions, this gives the modified index policy a complete practical story across n.

8. Conclusion

This paper studies how a decision-maker should prioritize and potentially abandon components of a multi-component project when completion of all components is required for full success and the feasibility of each component is uncertain. The key challenge in such environments is that stopping and sequencing cannot be separated: whether it is optimal to continue working on one component depends on the value of the remaining ones, while the choice of which component to pursue first depends on how quickly unsuccessful experimentation resolves uncertainty about the project as a whole.

Our first main result shows that, the dynamic control problem admits a sharp simplification. Although the decision-maker can in principle divide effort across components, there is no loss of optimality in reducing the problem to a choice of how long to work on each component. We then show, that in the feasibility-risk model, the optimal policy is a corner solution: the decision maker focuses entirely on one component until either that component is completed or the project is abandoned. The economic intuition is that if the project is ultimately infeasible, the decision-maker prefers to learn that as quickly as possible.

Our second set of results characterizes which component should be prioritized. These results show that the familiar intuition to "do the monkey first" is not universally correct. While lower prior probability makes a component more attractive to test first, the optimal ordering also depends on arrival rates and size of the reward. In particular, for sufficiently high rewards, the optimal policy takes a simple lexicographic form – begin with the component with faster arrival rate. In that sense, moonshot projects can rationally justify doing the pedestal first.

Because exact optimization becomes less transparent as projects grow in complexity, we also develop an approximate index policy that depends only on the characteristics of an individual component. The index has a natural economic interpretation in terms of failure risk relative to expected cost, performs well in the baseline environment, and becomes asymptotically optimal as rewards grow large. It therefore provides a tractable guide to prioritization even when solving the full dynamic problem is impractical.

Finally, the extension to general information structures clarifies the scope of the baseline results. The corner-solution property is not a generic feature of all learning problems. Rather, it is driven by uncertainty about feasibility itself. When uncertainty instead concerns the speed or cost of completion conditional on feasibility, interior solutions may arise and switching between components can become optimal. More broadly, this shows that the optimal organization of experimentation in multi-component projects depends critically on the nature of the underlying uncertainty.

Taken together, the results provide a tractable model of task prioritization and optimal stopping in complex projects with multiple necessary parts. They suggest that the conventional managerial advice to tackle the hardest task first should be applied with care: whether that intuition is correct depends not only on which component appears more feasible, but also on how fast it can be completed if feasible.

The analysis therefore refines the familiar advice to do the monkey before the pedestal. With a common opportunity cost of time, the weak-link logic is correct when components differ mainly in feasibility. But in dynamic moonshot projects, speed of learning also matters. When the reward is sufficiently large, the optimal first step is the component that resolves uncertainty fastest; if arrival rates are tied, the weak-link logic reappears and the less feasible component is tested first.

References

Agastya, Murali, Oleksii Birulin. 2023. Optimal task scheduling under adverse selection and hidden actions. *American Economic Journal: Microeconomics* **15**(2) 660–698. doi:10.1257/mic.20210126.

Alizamir, Saed, Francis de Véricourt, Peng Sun. 2022. Search under accumulated pressure. *Operations Research* **70**(3) 1393–1409. doi:10.1287/opre.2019.1880.

Babineaux, Ryan, John Krumboltz. 2013. *Fail fast, fail often: How losing can help you win*. Penguin.

Barker, B. 2023. Which comes first: The monkey or the pedestal? Deloitte, Insights to Action, Published online at `https://action.deloitte.com/insight/3380/which-comes-first-the-monkey-or-the-pedestal`, Accessed 7/6/2025.

Beyhaghi, Hedyeh, Linda Cai. 2024. Recent developments in pandora's box problem: Variants and applications. *ACM SIGecom Exchanges* **21**(1) 20–34.

Bickel, J. Eric, James E. Smith. 2006. Optimal sequential exploration: A binary learning model. *Decision Analysis* **3**(1) 16–32. doi:10.1287/deca.1050.0052.

Bickel, J. Eric, James E. Smith, Jennifer L. Meyer. 2008. Modeling dependence among geologic risks in sequential exploration decisions. *SPE Reservoir Evaluation & Engineering* **11**(2) 352–361. doi:10.2118/102369-PA.

Brown, D., J. Smith. 2013. Optimal sequential exploration: Bandits, clairvoyants, and wildcats. *Operations Research* **61**(3) 644–665.

Brown, David B., James E. Smith. 2020. Index policies and performance bounds for dynamic selection problems. *Management Science* **66**(7) 3029–3050. doi:10.1287/mnsc.2019.3342.

Carnehl, Christoph, Marco Ottaviani, Justus Preusser. 2025. Designing scientific grants. *Entrepreneurship and Innovation Policy and the Economy* **4** 139–178. doi:10.1086/732856.

Carnehl, Christoph, Johannes Schneider. 2023. On risk and time pressure: When to think and when to do. *Journal of the European Economic Association* **21**(1) 1–47.

Carnehl, Christoph, Johannes Schneider. 2025. A quest for knowledge. *Econometrica* **93**(2) 623–659. doi:10.3982/ECTA22144.

Chaimanowong, Wee, T. Tony Ke, J. Miguel Villas-Boas. 2023. Searching for breakthroughs. Available at SSRN 4425117. doi:10.2139/ssrn.4425117. Working paper.

Creemers, Stefan, Bert De Reyck, Roel Leus. 2015. Project planning with alternative technologies in uncertain environments. *European Journal of Operational Research* **242**(2) 465–476. doi:10.1016/j.ejor.2014.11.014.

Curello, Gregorio, Ludvig Sinander. 2025. Screening for breakthroughs. *Theoretical Economics* **20**(4) 1323–1365. doi:10.3982/TE6247.

De Reyck, Bert, Roel Leus. 2008. R&d project scheduling when activities may fail. *IIE Transactions* **40**(4) 367–384. doi:10.1080/07408170701413944.

Duke, A. 2022. Monkeys and pedestals: Find the bottleneck and solve for that first. Published online at `https://www.annieduke.com/newsletter-monkeys-and-pedestals-find-the-bottleneck-and-solve-for-that-first/`, Accessed 7/6/2025.

Erat, Sanjiv, Stylianos Kavadias. 2008. Sequential testing of product designs: Implications for learning. *Management Science* **54**(5) 956–968. doi:10.1287/mnsc.1070.0784.

Escobar, Juan F., Qiaoxi Zhang. 2021. Delegating learning. *Theoretical Economics* **16**(2) 571–603. doi:10.3982/TE4191.

Feng, Felix Zhiyu, Curtis R. Taylor, Mark M. Westerfield, Feifan Zhang. 2024. Setbacks, shutdowns, and overruns. *Econometrica* **92**(3) 815–847. doi:10.3982/ECTA21548.

Fleming, Lee. 2001. Recombinant uncertainty in technological search. *Management Science* **47**(1) 117–132. doi:10.1287/mnsc.47.1.117.10671.

Gittins, John C. 1979. Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society Series B: Statistical Methodology* **41**(2) 148–164.

Goldstein, Anna P., Michael Kearney. 2020. Know when to fold 'em: An empirical description of risk management in public research funding. *Research Policy* **49**(1) 103873. doi:10.1016/j.respol.2019.103873.

Green, Brett, Curtis R. Taylor. 2016. Breakthroughs, deadlines, and self-reported progress: Contracting for multistage projects. *American Economic Review* **106**(12) 3660–3699. doi:10.1257/aer.20151181.

Guo, Yingni. 2016. Dynamic delegation of experimentation. *American Economic Review* **106**(8) 1969–2008. doi:10.1257/aer.20141215.

Kahl, A. 2023. The monkey and the pedestal. Published online at `https://thebootstrappedfounder.com/the-monkey-and-the-pedestal/`, Accessed 7/6/2025.

Keller, Godfrey, Sven Rady. 2010. Strategic experimentation with poisson bandits. *Theoretical Economics* **5**(2) 275–311.

Keller, Godfrey, Sven Rady, Martin Cripps. 2005. Strategic experimentation with exponential bandits. *Econometrica* **73**(1) 39–68. doi:10.1111/j.1468-0262.2005.00564.x.

Kirpalani, Rishabh, Erik Madsen. 2023. Strategic investment evaluation. *Theoretical Economics* **18**(3) 1141–1180. doi:10.3982/TE4806.

Kwon, H. Dharma. 2010. Invest or exit? optimal decisions in the face of a declining profit stream. *Operations Research* **58**(3) 638–649. doi:10.1287/opre.1090.0740.

Kwon, H. Dharma, Steven A. Lippman. 2011. Acquisition of project-specific assets with bayesian updating. *Operations Research* **59**(5) 1119–1130. doi:10.1287/opre.1110.0949.

Kwon, H. Dharma, Steven A. Lippman, Christopher S. Tang. 2011. Sourcing decisions of project tasks with exponential completion times: Impact on operating profits. *International Journal of Production Economics* **134**(1) 138–150. doi:10.1016/j.ijpe.2011.06.007.

Madsen, Erik. 2022. Designing deadlines. *American Economic Review* **112**(3) 963–997. doi:10.1257/aer.20200212.

McCardle, Kevin F, Ilia Tsetlin, Robert L Winkler. 2018. When to abandon a research project and search for a new one. *Operations Research* **66**(3) 799–813.

McGrath, Rita Gunther. 2011. Failing by design. *Harvard Business Review* **89**(4) 76–83.

Moroni, Sofia. 2022. Experimentation in organizations. *Theoretical Economics* **17**(3) 1403–1450. doi:10.3982/TE3069.

Moscarini, Giuseppe, Lones Smith. 2001. The optimal level of experimentation. *Econometrica* **69**(6) 1629–1644. doi:10.1111/1468-0262.00259.

Oliveros, Santiago. 2025. Active product development. arXiv preprint arXiv:2507.15438.

Roberts, Kevin, Martin L. Weitzman. 1981. Funding criteria for research, development, and exploration projects. *Econometrica* **49**(5) 1261–1288.

Schmidt, Carl W., Ignacio E. Grossmann. 1996. Optimization models for the scheduling of testing tasks in new product development. *Industrial & Engineering Chemistry Research* **35**(10) 3498–3510. doi:10.1021/ie9601099.

Sitkin, Sim B. 1992. Learning through failure: The strategy of small losses. *Research in organizational behavior* **14** 231–266.

Teller, E. 2016. Tackle the monkey first. Published online at `https://blog.x.company/tackle-the-monkey-first-90fd6223e04d`, Accessed 7/6/2025.

Thijssen, Jacco Johan Jacob, Daniele Bregantini. 2017. Costly sequential experimentation and project valuation with an application to health technology assessment. *Journal of Economic Dynamics and Control* **77** 202–229. doi:10.1016/j.jedc.2017.01.016.

Thomke, Stefan H., David E. Bell. 2001. Sequential testing in product development. *Management Science* **47**(2) 308–323. doi:10.1287/mnsc.47.2.308.9838.

Urgun, Can, Leeat Yariv. 2025. Contiguous search: Exploration and ambition on uncharted terrain. *Journal of Political Economy* **133**(2) 522–567. doi:10.1086/732889.

Weitzman, M. L. 1979. Optimal search for the best alternative. *Econometrica* **47**(3) 641–54.

Zhao, Wenhui, Nicholas G. Hall, Zhixin Liu. 2020. Project evaluation and selection with task failures. *Production and Operations Management* **29**(2) 428–446. doi:10.1111/poms.13107.

EC.1. Derivations of some mathematical expressions.

Derivation of (13). Consider the strategy where we work on component 1 first for time T_1, and in the absence of success until T_1, we then work on component 2 for T_2 time, after which we stop if there is no success. If one of these components is completed, we switch to the optimal strategy for the remaining component, as given by (7). Denote the value function of a DM that follows such a strategy by $\tilde{V}(T_1,T_2)$. For a given realization of Λ_1 and Λ_2, the DM's payoff is then

$$\begin{aligned}\tilde{V}(T_1,T_2|\Lambda_1,\Lambda_2) = &\left(\frac{\Lambda_1}{\Lambda_1+\delta}\left(\frac{\Lambda_2 R-c}{\Lambda_2+\delta}\right)\left(1-e^{-(\Lambda_2+\delta)T_2^*}\right)-\frac{c}{\Lambda_1+\delta}\right)\left(1-e^{-(\Lambda_1+\delta)T_1}\right)\\ &+e^{-(\Lambda_1+\delta)T_1}\left(\frac{\Lambda_2}{\Lambda_2+\delta}\left(\frac{\Lambda_1 R-c}{\Lambda_1+\delta}\right)\left(1-e^{-(\Lambda_1+\delta)\left(T_1^*-T_1\right)}\right)-\frac{c}{\Lambda_2+\delta}\right)\left(1-e^{-(\Lambda_2+\delta)T_2}\right).\end{aligned} \tag{EC.1}$$

The expressions in the last two parentheses can be rewritten as

$$\begin{aligned}&\left(\frac{\Lambda_2}{\Lambda_2+\delta}\left(\frac{\Lambda_1 R-c}{\Lambda_1+\delta}\right)\left(e^{-(\Lambda_1+\delta)T_1}-e^{-(\Lambda_1+\delta)T_1^*}\right)-\frac{c}{\Lambda_2+\delta}e^{-(\Lambda_1+\delta)T_1}\right)\left(1-e^{-(\Lambda_2+\delta)T_2}\right)\\ &=\left(\frac{\Lambda_2}{\Lambda_2+\delta}\left(\frac{\Lambda_1 R-c}{\Lambda_1+\delta}\right)\left(e^{-(\Lambda_1+\delta)T_1}-1+1-e^{-(\Lambda_1+\delta)T_1^*}\right)-\frac{c}{\Lambda_2+\delta}\left(1-1+e^{-(\Lambda_1+\delta)T_1}\right)\right)\left(1-e^{-(\Lambda_2+\delta)T_2}\right)\\ &=\left(\frac{\Lambda_2}{\Lambda_2+\delta}\left(\frac{\Lambda_1 R-c}{\Lambda_1+\delta}\right)\left(1-e^{-(\Lambda_1+\delta)T_1^*}\right)-\frac{c}{\Lambda_2+\delta}\right)\left(1-e^{-(\Lambda_2+\delta)T_2}\right)\\ &+\left(\frac{\Lambda_2}{\Lambda_2+\delta}\left(\frac{\Lambda_1 R-c}{\Lambda_1+\delta}\right)\left(e^{-(\Lambda_1+\delta)T_1}-1\right)-\frac{c}{\Lambda_2+\delta}\left(e^{-(\Lambda_1+\delta)T_1}-1\right)\right)\left(1-e^{-(\Lambda_2+\delta)T_2}\right)\\ &=\left(\frac{\Lambda_2}{\Lambda_2+\delta}\left(\frac{\Lambda_1 R-c}{\Lambda_1+\delta}\right)\left(1-e^{-(\Lambda_1+\delta)T_1^*}\right)-\frac{c}{\Lambda_2+\delta}\right)\left(1-e^{-(\Lambda_2+\delta)T_2}\right)\\ &-\frac{\Lambda_1\Lambda_2 R-c\Lambda_2-c\Lambda_1-c\delta}{(\Lambda_1+\delta)(\Lambda_2+\delta)}\left(1-e^{-(\Lambda_1+\delta)T_1}\right)\left(1-e^{-(\Lambda_2+\delta)T_2}\right).\end{aligned}$$

Inserting this into the original expression (EC.1) and taking the expectation over Λ_1, Λ_2 yields

$$\begin{aligned}\tilde{V}(T_1,T_2) = &\,\mathbb{E}_{\Lambda_1}\left[\left(\frac{\Lambda_1}{\Lambda_1+\delta}V_2^*-\frac{c}{\Lambda_1+\delta}\right)\left(1-e^{-(\Lambda_1+\delta)T_1}\right)\right]\\ &+\mathbb{E}_{\Lambda_2}\left[\left(\frac{\Lambda_2}{\Lambda_2+\delta}V_1^*-\frac{c}{\Lambda_2+\delta}\right)\left(1-e^{-(\Lambda_2+\delta)T_2}\right)\right]\\ &-\mathbb{E}_{\Lambda_1,\Lambda_2}\left[\frac{\Lambda_1\Lambda_2 R-c\Lambda_2-c\Lambda_1-c\delta}{(\Lambda_1+\delta)(\Lambda_2+\delta)}\left(1-e^{-(\Lambda_1+\delta)T_1}\right)\left(1-e^{-(\Lambda_2+\delta)T_2}\right)\right].\end{aligned}$$

Notice that this expression is almost symmetrical: if we switch the order of components, it remains the same, thus yielding (13).

EC.2. Proofs of formal statements.

The following technical lemma ensures regularity properties π_i and $\hat{\Lambda}_i$ needed in several of the proofs that follow.

Lemma 1 (Properties of π_i and $\hat{\Lambda}_i$) *For each i, the functions π_i and $\hat{\Lambda}_i$ are continuously differentiable on $(0,\infty)$ with $\pi_i(0)=1$, $\lim_{t\to\infty}\pi_i(t)=\mathbb{P}(\Lambda_i=0)$ as $t\to\infty$, and $\hat{\Lambda}_i(0)=\mathbb{E}[\Lambda_i]$, $lim_{t\to\infty}\hat{\Lambda}_i(t)=\inf\operatorname{supp}(\Lambda_i)$. Moreover, π_i is log-convex and $\hat{\Lambda}_i$ is non-increasing.*

Proof. The log-convexity of π_i is the inequality $\pi_i(t)\,\mathbb{E}[\Lambda_i^2 e^{-\Lambda_i t}]\geq\mathbb{E}[\Lambda_i e^{-\Lambda_i t}]^2$, which is Cauchy–Schwarz inequality applied to the random variable Λ_i weighted by $e^{-\Lambda_i t}$. Differentiating $\hat{\Lambda}_i=-(\log\pi_i)'$ gives $\hat{\Lambda}_i'(t)=-(\log\pi_i)''(t)\leq 0$, so $\hat{\Lambda}_i$ is non-increasing. The limits at the endpoints are direct consequences of dominated convergence. □

Proof of Proposition 1. By Lemma 1, $\hat{\Lambda}_i$ is non-increasing and continuous and varies between $\hat{\Lambda}_i(0)=\mathbb{E}[\Lambda_i]$ and $\hat{\Lambda}_i(\infty)=\inf\operatorname{supp}(\Lambda_i)$. Combining with $V_i'(T_i)=e^{-\delta T_i}\pi_i(T_i)(R\hat{\Lambda}_i(T_i)-c)$ from (5) and $\pi_i>0$ yields (i)–(iii) and the uniqueness of the interior root in case (iii). □

Proof of Theorem 1. For each component i, let Y_i denote the (extended-real-valued) cumulative time on component i at which a success would occur if the DM worked on it continuously. Conditional on $\Lambda_i=\lambda$, $Y_i\sim\operatorname{Exp}(\lambda)$; marginalising over the prior yields $\mathbb{P}(Y_i>t)=\pi_i(t)$, with an atom $\mathbb{P}(Y_i=\infty)=\mathbb{P}(\Lambda_i=0)$ if the prior has a point mass on 0. The random variables Y_1,Y_2 are independent under the prior because the arrival processes on the two components are independent.

We first dispose of the trivial case $R\,\mathbb{E}[\Lambda_i]\leq c$ for some i: Proposition 1 then gives $T_i^*=0$, so the constraints $\int_0^T u(s)\,ds\leq T_1^*$ and $T-\int_0^T u(s)\,ds\leq T_2^*$ force $T=0$, and the claim is vacuous.

For the remainder, assume $R\,\mathbb{E}[\Lambda_i]>c$ for $i=1,2$. Fix any admissible $u\in\mathcal{U}$ and stopping time $T\in\mathbb{R}_+$ with $T_1:=\int_0^T u(s)\,ds\leq T_1^*$ and $T_2:=T-T_1\leq T_2^*$, and let $\hat{u}\in\mathcal{U}$ be any second admissible control with the same cumulative split, $\int_0^T\hat{u}(s)\,ds=T_1$.

The key path-independence result. Conditional on the realization (Y_1,Y_2), the project's discounted payoff depends only on (Y_1,Y_2) and (T_1,T_2), not on the particular control used to accumulate T_1 on component 1 and T_2 on component 2. We establish this in four cases on the realization (Y_1,Y_2) relative to (T_1,T_2).

Case 1: $Y_1>T_1$ and $Y_2>T_2$. Under any control with cumulative times (T_1,T_2), cumulative time on component i never exceeds $T_i<Y_i$, so neither component succeeds during $[0,T]$. The DM accumulates cost at rate c for the full horizon T and receives no reward. The realized discounted payoff equals $-c\,(1-e^{-\delta T})/\delta$, which depends only on $T=T_1+T_2$ and not on the path.

Case 2: $Y_1\leq T_1$ and $Y_2\leq T_2$. Both components eventually succeed. Let τ denote the calendar time of the first success—i.e., the earliest t at which $t_1(t;u)=Y_1$ or $t_2(t;u)=Y_2$, where $t_i(t;u)$ denotes cumulative time on component i at calendar time t under control u. Say, without loss of generality, $t_1(\tau;u)=Y_1$ (the symmetric case $t_2(\tau;u)=Y_2$ is handled identically). At calendar time τ the cumulative time on component 2 is $t_2(\tau;u)=\tau-Y_1$.

After this first success, the DM switches to the single-component optimal continuation on component 2 (eq. (7)), which prescribes working continuously on component 2 until either success or cumulative time on 2 reaches T_2^*. Since $Y_2 \le T_2 \le T_2^*$, the success threshold is reached first. The additional working time needed on component 2 is $Y_2-(\tau-Y_1)=Y_1+Y_2-\tau$, so the second success occurs at calendar time $\tau+(Y_1+Y_2-\tau)=Y_1+Y_2$, which is independent of both τ and the path.

The DM works continuously from time 0 until the project completes at Y_1+Y_2, so the discounted cost is $c(1-e^{-\delta(Y_1+Y_2)})/\delta$ and the discounted reward is $Re^{-\delta(Y_1+Y_2)}$. Both depend only on (Y_1,Y_2) and not on the path.

Case 3: $Y_1 \le T_1$ **and** $Y_2 > T_2$**.** Under any admissible control with cumulative times (T_1,T_2), component 1's success requirement Y_1 is reached within $[0,T]$ (because cumulative time on 1 increases continuously from 0 to $T_1 \ge Y_1$), whereas component 2's requirement Y_2 is not reached during $[0,T]$ (because cumulative time on 2 is bounded above by $T_2 < Y_2$). Hence the only success during $[0,T]$ is on component 1, at some calendar time $\tau_1 \in [Y_1,T]$ depending on the path.

After success on 1 at τ_1, the DM switches to single-component optimal continuation on component 2, starting from cumulative time $t_2(\tau_1;u)=\tau_1-Y_1$ on 2. The continuation works on component 2 until cumulative time on 2 reaches either Y_2 (success, additional time $Y_2-(\tau_1-Y_1)$) or T_2^* (abandonment, additional time $T_2^*-(\tau_1-Y_1)$), whichever comes first.

- Sub-case 3a: $Y_2 \le T_2^*$. Success on 2 occurs first. Second success at calendar time $\tau_1+(Y_2-(\tau_1-Y_1))=Y_1+Y_2$. Discounted reward $Re^{-\delta(Y_1+Y_2)}$ and discounted cost $c(1-e^{-\delta(Y_1+Y_2)})/\delta$.
- Sub-case 3b: $Y_2 > T_2^*$. Abandonment first. Project ends at calendar time $\tau_1+(T_2^*-(\tau_1-Y_1))=Y_1+T_2^*$. Discounted reward 0 and discounted cost $c(1-e^{-\delta(Y_1+T_2^*)})/\delta$.

In either sub-case, the relevant calendar times collapse to functions of (Y_1,Y_2) and T_2^* alone; the path-dependent τ_1 cancels. The discounted payoff depends only on (Y_1,Y_2).

Case 4: $Y_1 > T_1$ **and** $Y_2 \le T_2$**.** Symmetric to Case 3 with the roles of components 1 and 2 exchanged.

Conclusion. The realized discounted payoff is in every case a function of $(Y_1,Y_2,T_1,T_2,T_1^*,T_2^*)$ alone— the joint law of (Y_1,Y_2) depends only on the prior on (Λ_1,Λ_2), and the four quantities T_1,T_2,T_1^*,T_2^* are determined by the cumulative split and the model primitives. Hence the expected discounted payoff under u and under $\hat{u}$ are equal, completing the proof of payoff-equivalence at fixed (T_1,T_2).

It remains to verify that any optimal u satisfies $\int_0^T u(s)\,ds \le T_1^*$ and $T-\int_0^T u(s)\,ds \le T_2^*$. This is immediate from the single-component analysis: the expected reward for *completing* component i

in any continuation is strictly less than R (since $V_j^* < R$ for $j \neq i$), so the single-component abandonment rule of Proposition 1 prescribes $T_i \leq T_i^*$ for each i at any optimum. $\square$

The lemma below is a technical one which is used in the proof of Theorem 2. Both the lemma and the theorem that follow use the notation $\Lambda_i(t_i)$ which stands for the posterior for Λ_i after working on component i for cumulative time t_i with no success.

Lemma 2 $\mathrm{Var}(\Lambda_i(t_i)) = -d\hat{\Lambda}_i(t_i)/dt_i$.

Proof. Conditional on no success before time t_i, from Bayes' theorem we have

$$\hat{\Lambda}_i(t_i) = \frac{\int_0^\infty \lambda e^{-\lambda t_i} dF_i(\lambda)}{\int_0^\infty e^{-\lambda t_i} dF_i(\lambda)},$$

$$\begin{aligned}\frac{d\hat{\Lambda}_i(t_i)}{dt_i} &= -\frac{\int_0^\infty \lambda^2 e^{-\lambda t_i} dF_i(\lambda)}{\int_0^\infty e^{-\lambda t_i} dF_i(\lambda)} + \frac{\left(\int_0^\infty \lambda e^{-\lambda t_i} dF_i(\lambda)\right)^2}{\left(\int_0^\infty e^{-\lambda t_i} dF_i(\lambda)\right)^2} \\ &= -\mathbb{E}[\Lambda_i^2(t_i)] + \hat{\Lambda}_i(t_i)^2 = -\mathrm{Var}\left(\Lambda_i(t_i)\right). \qquad \square\end{aligned}$$

Proof of Theorem 2. Assume that $V(T_1, T_2)$ is maximized at an interior point. Then, since the objective function of (14) is twice differentiable, there must exist a prior over arrival rates such that the first-order conditions (FOC) and second-order conditions (SOC) for a maximum are satisfied at $T_1 = T_2 = 0$. In particular, the gradient must be zero and the Hessian must be negative semi-definite. We show when such a condition is violated.

To emphasize the sequential structure of the decision problem, we rewrite (13) as:

$$\begin{aligned}V(T_1, T_2) = {} & \mathbb{E}_{\Lambda_1}\left[\left(\frac{\Lambda_1}{\Lambda_1+\delta}V_2^* - \frac{c}{\Lambda_1+\delta}\right)\left(1 - e^{-(\Lambda_1+\delta)T_1}\right)\right] \\ & + \mathbb{E}_{\Lambda_1,\Lambda_2}\left[e^{-(\Lambda_1+\delta)T_1}\left(\frac{\Lambda_2}{\Lambda_2+\delta}\left(\frac{\Lambda_1 R - c}{\Lambda_1+\delta}\right)\left(1 - e^{-(\Lambda_1+\delta)(T_1^*-T_1)}\right) - \frac{c}{\Lambda_2+\delta}\right)\left(1 - e^{-(\Lambda_2+\delta)T_2}\right)\right].\end{aligned} \tag{EC.2}$$

This form highlights the sequential structure: the decision-maker works on component 1 first, and, conditional on success, proceeds to component 2. The second term in (EC.2) can be factored as a product of expectations:

$$\mathbb{E}_{\Lambda_1}\left[e^{-(\Lambda_1+\delta)T_1}\right] \cdot \mathbb{E}_{\Lambda_2}\left[\left(\frac{\Lambda_2}{\Lambda_2+\delta}V_1^*(T_1) - \frac{c}{\Lambda_2+\delta}\right)\left(1 - e^{-(\Lambda_2+\delta)T_2}\right)\right],$$

where

$$V_1^*(T_1) = \frac{\mathbb{E}_{\Lambda_1}\left[e^{-(\Lambda_1+\delta)T_1}\left(\frac{\Lambda_1 R - c}{\Lambda_1+\delta}\right)\left(1 - e^{-(\Lambda_1+\delta)(T_1^*-T_1)}\right)\right]}{\mathbb{E}_{\Lambda_1}\left[e^{-(\Lambda_1+\delta)T_1}\right]} \tag{EC.3}$$

denotes the expected value from component 1, updated after working on it for time T_1. At $T_1 = 0$, we have $V_1^*(0) = V_1^*$.

Differentiating $V(T_1,T_2)$ and applying the simplification from (4), we obtain the following derivatives:

$$\frac{\partial V(0,T_2)}{\partial T_2} = \mathbb{E}_{\Lambda_2}\left[(V_1^*\Lambda_2 - c)\, e^{-(\Lambda_2+\delta)T_2}\right], \tag{EC.4}$$

$$\frac{\partial V(0,0)}{\partial T_1} = \mathbb{E}_{\Lambda_1}\left[V_2^*\Lambda_1 - c\right] = 0, \tag{EC.5}$$

$$\frac{\partial V(0,0)}{\partial T_2} = \mathbb{E}_{\Lambda_2}\left[V_1^*\Lambda_2 - c\right] = 0. \tag{EC.6}$$

Equations (EC.5) and (EC.6) correspond to the first-order conditions at $(T_1,T_2) = (0,0)$. We will next analyze the second-order conditions (SOC) to determine when a maximum at the interior is ruled out.

To construct the Hessian, we differentiate again:

$$\frac{\partial^2 V(0,0)}{\partial T_2^2} = -V_1^*\,\mathrm{Var}(\Lambda_2), \tag{EC.7}$$

$$\frac{\partial^2 V(0,0)}{\partial T_1^2} = -V_2^*\,\mathrm{Var}(\Lambda_1), \tag{EC.8}$$

$$\frac{\partial^2 V(0,0)}{\partial T_1 \partial T_2} = \hat{\Lambda}_2 \cdot \frac{dV_1^*}{dT_1} = \hat{\Lambda}_1 \cdot \frac{dV_2^*}{dT_2}. \tag{EC.9}$$

Putting it together, the determinant of the Hessian becomes:

$$H = V_1^* V_2^*\,\mathrm{Var}(\Lambda_1)\,\mathrm{Var}(\Lambda_2) - \hat{\Lambda}_1\hat{\Lambda}_2 \cdot \frac{dV_1^*}{dT_1} \cdot \frac{dV_2^*}{dT_2}.$$

Using the identity $\mathrm{Var}(\Lambda_i) = -\frac{d\hat{\Lambda}_i}{dT_i}$ proved in Lemma 2 in Appendix A, we get:

$$H = V_1^* V_2^* \frac{d\hat{\Lambda}_1}{dT_1}\frac{d\hat{\Lambda}_2}{dT_2} - \hat{\Lambda}_1\hat{\Lambda}_2 \cdot \frac{dV_1^*}{dT_1} \cdot \frac{dV_2^*}{dT_2}.$$

So, $H < 0$ (i.e., the Hessian is indefinite) when $V_i^*(t_i)/\hat{\Lambda}_i(t_i)$ is decreasing in t_i. □

Proof of Proposition 2. By Theorem 2, it suffices to show that $V_i^*(t_i)/\hat{\Lambda}_i(t_i)$ is strictly decreasing in t_i. We do this for a generic continuation reward $\rho > 0$, since the argument makes no use of the specific value of the reward beyond positivity; we then specialize to the reward that appears in the two-component value function (13).

For any continuation reward $\rho > 0$ and any prior probability $p \in (0,1)$ of feasibility, the single-component value at prior p, arrival rate λ_i, cost c, and discount rate δ is, by (3)–(7),

$$V_i^*(p;\rho) = \frac{p(\lambda_i\rho - c)}{\lambda_i + \delta}\left(1 - e^{-(\lambda_i+\delta)\,T_i^*(p;\rho)}\right) - \frac{(1-p)\,c}{\delta}\left(1 - e^{-\delta\,T_i^*(p;\rho)}\right), \tag{EC.10}$$

where $T_i^*(p;\rho)$ is the optimal stopping time from (7),

$$T_i^*(p;\rho) = \frac{1}{\lambda_i}\log\frac{p(\lambda_i\rho - c)}{c(1-p)}.$$

For the two-point distribution, the posterior expected rate is $\hat{\Lambda}_i = p\,\lambda_i$, so

$$\frac{V_i^*(p;\rho)}{\hat{\Lambda}_i} = \frac{V_i^*(p;\rho)}{p\,\lambda_i} = \frac{\lambda_i\rho - c}{\lambda_i(\lambda_i+\delta)}\Big(1 - e^{-(\lambda_i+\delta)\,T_i^*(p;\rho)}\Big) - \frac{(1-p)\,c}{p\,\lambda_i\,\delta}\Big(1 - e^{-\delta\,T_i^*(p;\rho)}\Big). \qquad \text{(EC.11)}$$

Differentiating (EC.11) with respect to p and applying the envelope theorem (so the terms involving $\partial T_i^*/\partial p$ vanish) yields

$$\frac{d}{dp}\left(\frac{V_i^*(p;\rho)}{p\,\lambda_i}\right) = \frac{c\left(1 - e^{-\delta\,T_i^*(p;\rho)}\right)}{p^2\,\lambda_i\,\delta} \geq 0, \qquad \text{(EC.12)}$$

with equality only at the boundary $T_i^*(p;\rho) = 0$ (i.e., the project is not started). Hence the ratio $V_i^*(p;\rho)/\hat{\Lambda}_i$ is strictly increasing in p on the interior of the active region $\{p : p\,\lambda_i\rho > c\}$, for *every* continuation reward $\rho > 0$.

By Bayesian updating in the absence of success, the posterior probability of feasibility $p_i(t_i)$ is strictly decreasing in t_i. Composing with (EC.12) therefore gives that $V_i^*\big(p_i(t_i);\rho\big)/\hat{\Lambda}_i(t_i)$ is strictly decreasing in t_i for every $\rho > 0$.

In (14), the continuation reward that appears in the single-component sub-problem upon success on component j is precisely $\rho = V_j^*$, the single-component optimum at the prior. Setting $\rho = V_j^*$ in the conclusion above, the ratio $V_i^*(t_i)/\hat{\Lambda}_i(t_i)$ is strictly decreasing in t_i, so the hypothesis of Theorem 2 is satisfied and (14) admits no interior solution. $\square$

Proof of Corollary 1 From (15), $\lim_{\lambda_i\to\infty} V_i^{**} = p_i V_j^*(p_j) > \lim_{\lambda_i\to\infty} V_j^{**}$, as $\lim_{\lambda_i\to\infty} V_j^{**}$ is equal to $V_j^*(p_j)$, just with R replaced by $p_i R$. The statement of the corollary then follows from continuity of V_i^{**} in λ_i. $\square$

Proof of Corollary 2 Assume the opposite: $\exists p_1, p_2, \lambda_1, \lambda_2$ s. t. it is optimal to work on component 1 first, and $\tilde{p}_1 < p_1$ such that under $\tilde{p}_1, p_2, \lambda_1, \lambda_2$ it is optimal to work on component 2 first. But then, there exists t such that $p_i(t) = \tilde{p}_i$, so it will be optimal for the DM to work on component 1 for time t and then switch to component 2, which is in contradiction to the "no interior solutions" result of Proposition 2. $\square$

Proof of Proposition 3. Consider the Taylor expansion of (15) around $1/R = 0$, which yields

$$V_1^{**} = \frac{p_1 p_2 \lambda_1 \lambda_2 R}{(\lambda_1+\delta)(\lambda_2+\delta)} - \frac{c p_1 \lambda_1}{\lambda_1+\delta}\left(\frac{p_2}{\lambda_2+\delta} + \frac{1-p_2}{\delta}\right) - c\left(\frac{p_1}{\lambda_1+\delta} + \frac{1-p_1}{\delta}\right) + O\left(\frac{1}{R}\right).$$

Comparing this expression for V_1^{**} to the one for V_2^{**} yields that on some open set around $1/R = 0$, it is optimal to do first the component for which the following index is higher:

$$c\left(1 - \frac{p_i\lambda_i}{\lambda_i+\delta}\right)\left(\frac{p_j}{\lambda_j+\delta} + \frac{1-p_j}{\delta}\right).$$

Expanding and reorganizing this expression yields that it is equal to

$$c\frac{(\delta+\lambda_1(1-p_1))(\delta+\lambda_2(1-p_2))}{\delta(\lambda_1+\delta)(\lambda_2+\delta)}.$$

However, this expression is symmetric in i and j, thus the same for both components.

Thus, we must examine higher-order terms of the Taylor expansion to detect any asymmetries. Among the remaining terms, those that involve higher values of λ_i are more influential. Considering the second-order Taylor expansion, when $\lambda_1 \geq \lambda_2$, the comparison reduces to:

$$\frac{(1-p_1)\lambda_1}{\delta(\lambda_1+\delta)}\left(\frac{c(1-p_1)}{p_1(V_j^*\lambda_1-c)}\right)^{\delta/\lambda_1} \sim \frac{p_2\lambda_2(1-p_1)\lambda_1}{(\lambda_2+\delta)\delta(\lambda_1+\delta)}\left(\frac{c(1-p_1)}{p_1(R\lambda_1-c)}\right)^{\delta/\lambda_1}.$$

Here, left-hand side dominates (unless $\lambda_1=\lambda_2$ in which case one with lower p_i does), establishing the lexicographic policy of the proposition. □

Proof of Proposition 4. From (3), we can express $\hat{c}_i$ as

$$\begin{aligned}
\hat{c}_i &:= \mathbb{E}_{\Lambda_i}\left[\int_0^{T_i^*}\Lambda_i e^{-\Lambda_i t}c\frac{1-e^{-t\delta}}{\delta}dt + e^{-\Lambda_i T_i^*}c\frac{1-e^{-T_i^*\delta}}{\delta}\right] \\
&= \mathbb{E}_{\Lambda_i}\left[\frac{c\left(\delta+\Lambda_i e^{-T_i^*(\delta+\Lambda_i)}-(\delta+\Lambda_i)e^{-T_i^*\Lambda_i}\right)}{\delta(\delta+\Lambda_i)} + e^{-\Lambda_i T_i^*}c\frac{1-e^{-T_i^*\delta}}{\delta}\right] \\
&= p_i\frac{c\left(\delta+\lambda_i e^{-T_i^*(\delta+\lambda_i)}-(\delta+\lambda_i)e^{-T_i^*\lambda_i}\right)}{\delta(\delta+\lambda_i)} + (p_i e^{-\lambda_i T_i^*}+1-p_i)c\frac{1-e^{-T_i^*\delta}}{\delta}.
\end{aligned}$$

Similarly, we can write the expected reward of such agent (who receives reward r when completing the component) as

$$\begin{aligned}
\hat{r}_i r &:= r\mathbb{E}_{\Lambda_i}\left[\int_0^{T_i^*}\Lambda_i e^{-(\Lambda_i+\delta)t}dt\right] \\
&= r\mathbb{E}_{\Lambda_i}\left[\frac{\Lambda_i(1-e^{-T_i^*(\delta+\Lambda_i)})}{\delta+\Lambda_i}\right] \\
&= rp_i\frac{\lambda_i(1-e^{-T_i^*(\delta+\lambda_i)})}{\delta+\lambda_i}.
\end{aligned}$$

With this in mind we can write

$$\begin{aligned}
V(T_1^*,0) &= -\hat{c}_1+\hat{r}_1(\hat{r}_2R-\hat{c}_2), \\
V(0,T_2^*) &= -\hat{c}_2+\hat{r}_2(\hat{r}_1R-\hat{c}_1).
\end{aligned}$$

Thus, $V(T_1^*,0)\geq V(0,T_2^*)$ is equivalent to

$$\begin{aligned}
-\hat{c}_1-\hat{r}_1\hat{c}_2 &\geq -\hat{c}_2-\hat{r}_2\hat{c}_1, \\
\frac{1-\hat{r}_1}{\hat{c}_1} &\geq \frac{1-\hat{r}_2}{\hat{c}_2}.
\end{aligned}$$

Inserting back the expressions for $\hat{c}_i$ and $\hat{r}_i$ and simplifying the result yields (18). $\square$

Proof of Proposition 5.

Part 1: Deriving (19). Consider the situation where component i (the one with higher index) is infeasible; this occurs with probability $1-p_i$. In this case, the optimal policy will work on this component until T_i^{**} before giving up, while the index policy will work until T_i^*. Since $T_i^* > T_i^{**}$, that means that the index policy will incur additional costs between T_i^* and T_i^{**}. The present value of those costs is

$$\int_{T_i^{**}}^{T_i^*} e^{-\delta t} c\,dt = c\frac{e^{-\delta T_i^{**}} - e^{-\delta T_i^*}}{\delta}. \tag{EC.13}$$

The same situation can also occur even if component i is feasible, as that can also result in no arrivals until T_i^*, which happens with probability $e^{-\lambda_i T_i^*}$. Thus, multiplying the additional costs of (EC.13) by the probability of that scenario occurring yields

$$c(1-p_i+p_i e^{-\lambda_i T_i^*})\frac{e^{-\delta T_i^{**}} - e^{-\delta T_i^*}}{\delta}, \tag{EC.14}$$

which is the second additive term in (19).

We still need to consider what happens if there is (a) an arrival before T_i^{**} or (b) an arrival between T_i^{**} and T_i^*.

If a success is encountered before T_i^{**}, both the optimal and index policy will proceed in the same way, switching to the other component and then working on it for V_j^* or until success, whichever comes first. Thus, there will no difference in payoffs in this scenario.

Thus, the last scenario to consider is the one where arrival occurs between T_i^{**} and T_i^*. Denoting this arrival time with with τ, this scenario will cause the index policy to incur additional costs compared to the optimal one (from additional work on component i from T_i^{**} to τ). The present value of those costs is $c(e^{-\delta T_i^{**}} - e^{-\delta\tau})/\delta$. However, this gives the index policy the opportunity to possibly complete the project in this scenario, yielding an expected payoff of $e^{-\delta\tau}V_j^*$. Combining this additional cost and possible payoff with the density of τ that is equal to $\lambda_i \exp\{-\lambda_i\tau\}$ if component i is feasible yields

$$p_i\int_{T_i^{**}}^{T_i^*} \lambda_i e^{-\lambda_i\tau}\left(\frac{e^{-\delta T_i^{**}} - e^{-\delta\tau}}{\delta}c - e^{-\delta\tau}V_j^*\right)d\tau.$$

Solving this integral and adding (EC.14) to it gives (19).

Part 2: Deriving (20). As in the proof of Theorem 1, we denote by the random variable Y_i (with realization y_i) the arrival time of component i, and separate the proof into several mutually and collectively exclusive cases depending on realizations of this variable. We additionally separate the cases depending on whether the index policy misorders the component prioritization compared to the optimal one. As in the first part of the proof, we assume $I_i \geq I_j$ without loss of generality.

Case 1: No misordering, $y_i \leq T_i^{**}$, $y_j \leq T_j^*$. Here both policies will yield exactly the same payoff, as they will both follow the same order and duration of work, and receive the reward at the same time.

Case 2: No misordering, $T_i^{**} < y_i$. Here, the optimal policy will stop working at T_i^{**} while the index policy will continue. The worst sample path for the index policy in this case is one where continues working for the maximum possible time without ultimately receiving the reward. This happens if it encounters success in component i at exactly the time where it would otherwise give up (T_i^*) and then works on component j for full time (T_j^*) without success, resulting in additional (discounted) costs of $c(e^{-\delta T_i^{**}} - e^{-\delta T_i^*})/\delta + c(e^{-\delta T_i^*} - e^{-\delta(T_i^*+T_j^*)})/\delta$.

Case 3: misordering, $y_j \leq T_j^{**}$, $y_i \leq T_i^*$. Here, both policies will receive the reward at exactly the same time and incur identical costs, rendering the misordering inconsequential.

Case 4: misordering, $T_j^{**} < y_j$. Analogously to Case 2, but with the addition of misordering, here the maximal cost difference is $c(-e^{-\delta(T_i^*+T_j^*)} + e^{-\delta T_j^{**}})/\delta$.

Noting that the costs in Case 4 are greater than those in Case 3 (intuitively, this is because the index policy in Case 4 experiences the same misordering costs as in Case 3, but also incurs a greater sum of total costs, which does not happen in Case 3) (20) is obtained by taking the maximum of Cases 2 and 4.

Proof of Proposition 6. First, we show that for sufficiently high R, the index policy prioritizes the same component as the fully optimal policy. Denoting $\rho = 1/R$, the asymptotic expansion of (18) around $\rho = 0$ yields the Puiseux series

$$I_i = \frac{\delta}{c} + \frac{(1-p_i)(\delta+\lambda_i)}{c\,\delta\big(\delta+\lambda_i(1-p_i)\big)}\left(\frac{c(1-p_i)}{p_i\lambda_i}\right)^{\delta/\lambda_i}\rho^{\delta/\lambda_i} + O\left(\rho^{1+\delta/\lambda_i}\right), \qquad \rho \to 0^+. \tag{EC.15}$$

The leading term δ/c is the same for both components, so the ranking of I_1 and I_2 for small ρ is determined by the subleading correction. The coefficient of ρ^{δ/λ_i} in (EC.15) is strictly positive; a higher arrival rate λ_i gives a smaller exponent δ/λ_i and hence a larger ρ^{δ/λ_i} factor (since $\rho < 1$ eventually). Consequently the index policy selects the component with the *higher* arrival rate for R sufficiently large.

If the arrival rates are equal, $\lambda_1 = \lambda_2 = \lambda$, the exponents on ρ coincide and the comparison reduces to that of the coefficients. Dropping the indices on the parameters that are shared, $I_i > I_j$ amounts to

$$\frac{1-p_i}{\delta+\lambda(1-p_i)}\left(\frac{1-p_i}{p_i}\right)^{\delta/\lambda} > \frac{1-p_j}{\delta+\lambda(1-p_j)}\left(\frac{1-p_j}{p_j}\right)^{\delta/\lambda}. \tag{EC.16}$$

The expression on each side, viewed as a function of the corresponding p, is strictly decreasing on $(0,1)$. Hence (EC.16) holds if and only if $p_i < p_j$: the index policy selects the component with the *lower* p_i. Combined with the λ-ranking above, the index thus implements the same lexicographic policy that is fully optimal for high R by Proposition 3.

With the correct-component selection established for $R > \bar{R}$, the remaining convergence statements follow from $\lim_{R\to\infty} T_i^* = \lim_{R\to\infty} T_i^{**} = \infty$ for $i = 1,2$: taking the limit as $T_1, T_2, R \to \infty$ in (13) gives $\lim_{R\to\infty} V(I) = \lim_{R\to\infty} \max_k V_k^{**}$. Finally, $\lim_{R\to\infty} B(I) = 0$ follows from the bound established in the proof of Proposition 5: when the index policy correctly orders the components, $B(I) \le c\big(e^{-\delta T_i^{**}} - e^{-\delta(T_i^* + T_j^*)}\big)/\delta$, and the right-hand side tends to 0 as $R \to \infty$. $\square$

Proof of Corollary 3. The statement follows directly from Proposition 6. $\square$

Proof of Proposition 7. By Theorem 2, it suffices to verify that $(V_i^*(t_i)/\hat{\Lambda}_i(t_i))' \le 0$ at any candidate interior critical point of (14) for all R sufficiently large. We prove this in five steps.

Step 1: Interior critical points sit at large times. The interior first-order conditions read $V_i^*(t_i)\,\hat{\Lambda}_j(t_j) = c,\ (i,j) \in \{(1,2),(2,1)\}$. The single-component value $V_i^*(t_i)$ is continuous in t_i, strictly positive on $[0, T_i^*)$, with $V_i^*(T_i^*) = 0$. Let $\hat{r}_i^\infty := \mathbb{E}\left[\Lambda_i/(\Lambda_i + \delta)\right]$, which is strictly positive because $\inf\big(\mathrm{supp}(\Lambda_i) \cap (0,\infty)\big) > 0$ by assumption. We claim $V_i^*(0)/R \to \hat{r}_i^\infty$ as $R \to \infty$. On the one hand, every fixed-horizon policy with time T has prior value at most $R\,\mathbb{E}\big[\frac{\Lambda_i(1-e^{-(\Lambda_i+\delta)T})}{\Lambda_i+\delta}\big] \le R\,\hat{r}_i^\infty$, so $V_i^*(0) \le R\,\hat{r}_i^\infty$ for every R. On the other hand, working for any fixed $T_0 > 0$ and then abandoning is a feasible single-component policy whose prior value

$$R\,\mathbb{E}\left[\tfrac{\Lambda_i(1-e^{-(\Lambda_i+\delta)T_0})}{\Lambda_i+\delta}\right] - c\,\mathbb{E}\left[\tfrac{1-e^{-(\Lambda_i+\delta)T_0}}{\Lambda_i+\delta}\right]$$

is a lower bound on $V_i^*(0)$, and the R-coefficient increases to $\hat{r}_i^\infty$ as $T_0 \to \infty$. In particular, $V_i^*(0) \ge \frac{1}{2}\hat{r}_i^\infty\,R$ for all R sufficiently large. The right-hand side $c/\hat{\Lambda}_j(t_j)$ of the FOC is bounded above by $c/\hat{\Lambda}_j(T_j^*) = c/(c/R) = R$, hence bounded. Combined with $V_i^*(0)/R \to \hat{r}_i^\infty > 0$, we conclude that the level set $\{t_i : V_i^*(t_i) \le c/\hat{\Lambda}_j(t_j)\}$ pushes to $T_i^* \to \infty$ as $R \to \infty$. In particular, $t_i \to \infty$ for any interior critical point.

Step 2: Posterior decomposition and asymptotic concentration. Decompose $\Lambda_i(t_i) = 0 \cdot \mathbf{1}\{\Lambda_i = 0\} + \Lambda_i\,\mathbf{1}\{\Lambda_i > 0\}$. The posterior puts mass $q_i(t_i) = p_{i,0}/\pi_i(t_i)$ at 0 and the rest on the positive support; let $\eta_i(t_i) := 1 - q_i(t_i)$. Because $\pi_i(t_i) \downarrow p_{i,0} > 0$, we have $\eta_i(t_i) \downarrow 0$ as $t_i \to \infty$.

Under the gap assumption $\lambda_{i1} > 0$, the posterior of Λ_i conditional on $\{\Lambda_i > 0\} \cap \{\tau_i > t_i\}$ concentrates on λ_{i1} as $t_i \to \infty$. Explicitly, for any $\varepsilon > 0$,

$$\mathbb{P}\big(\Lambda_i > \lambda_{i1} + \varepsilon \;\big|\; \Lambda_i > 0, \tau_i > t_i\big) = \frac{\displaystyle\int_{\lambda_{i1}+\varepsilon}^{\infty} e^{-\lambda t_i}\, dF_i^+(\lambda)}{\displaystyle\int_{\lambda_{i1}}^{\infty} e^{-\lambda t_i}\, dF_i^+(\lambda)} \le e^{-\varepsilon t_i} \to 0,$$

where F_i^+ is the prior restricted to the positive support. Therefore

$$\mu_i^+(t_i) \to \lambda_{i1}, \qquad \mu_i^{+(2)}(t_i) \to \lambda_{i1}^2, \qquad \mathrm{Var}_i^+(t_i) := \mu_i^{+(2)}(t_i) - \mu_i^+(t_i)^2 \to 0, \qquad t_i \to \infty.$$

Step 3: Closed-form for the slope at the interior FOC. At any candidate interior FOC point, $V_i^*(t_i) = c/\hat{\Lambda}_j(t_j) = c/(\eta_j \mu_j^+)$. Using the conditional values $V_{i,0}(t_i) := V_i^*(t_i \mid \Lambda_i = 0) = -c(1 - e^{-\delta(T_i^* - t_i)})/\delta \le 0$ and $V_{i,\lambda}(t_i) := V_i^*(t_i \mid \Lambda_i = \lambda)$ for $\lambda > 0$, we have $V_i^*(t_i) = q_i V_{i,0}(t_i) + \eta_i \mathbb{E}_{t_i}^+[V_{i,\Lambda_i}(t_i)]$, and

$$\mathrm{Cov}(\Lambda_i(t_i), V_i^*(t_i \mid \Lambda_i(t_i))) = \eta_i \Big[(1-\eta_i)\mu_i^+(t_i)\big(\mathbb{E}_{t_i}^+[V_{i,\Lambda_i}(t_i)] - V_{i,0}(t_i)\big) + \mathrm{cov}_i^+\Big],$$

where $\mathrm{cov}_i^+ := \mathbb{E}_{t_i}^+[\Lambda_i V_{i,\Lambda_i}(t_i)] - \mu_i^+(t_i)\, \mathbb{E}_{t_i}^+[V_{i,\Lambda_i}(t_i)] \ge 0$ by Chebyshev's sum inequality applied to the increasing functions $\lambda \mapsto \lambda$ and $\lambda \mapsto V_{i,\lambda}$.

The numerator of $(V_i^*/\hat{\Lambda}_i)'(t_i)$ is $\mathrm{Num}_i := V_i^* \,\mathrm{Var}(\Lambda_i) - \hat{\Lambda}_i \,\mathrm{Cov}_i$. Substituting $V_i^* = c/(\eta_j \mu_j^+)$, the variance decomposition $\mathrm{Var}(\Lambda_i) = \eta_i \mu_i^{+(2)} - \eta_i^2 (\mu_i^+)^2$, and the FOC identity $\eta_i \mu_i^+ \big[(1-\eta_i)V_{i,0} + \eta_i \mathbb{E}^+[V_{i,\Lambda_i}]\big] = c\mu_i^+/(\eta_j \mu_j^+)$, a straightforward computation (cf. the analogous calculation in the two-point proof of Proposition 2) yields

$$\mathrm{Num}_i(t_i, t_j) = \underbrace{\frac{c\,\mathrm{Var}_i^+(t_i)}{\eta_j \mu_j^+ \mu_i^+}}_{=:A(t_i,t_j)} + \underbrace{\eta_i (\mu_i^+)^2 V_{i,0}(t_i)}_{=:B(t_i)} - \underbrace{\eta_i^2 \mu_i^+ \,\mathrm{cov}_i^+}_{=:C(t_i)}. \tag{EC.17}$$

Here $A \ge 0$, $B \le 0$ (since $V_{i,0} \le 0$), and $C \ge 0$.

Step 4: Asymptotic dominance of B over A. By Step 1, the candidate $t_i \to \infty$. By Step 2, $\mathrm{Var}_i^+(t_i) \to 0$ as $t_i \to \infty$. Combining with $\mu_i^+(t_i) \to \lambda_{i1} > 0$,

$$A(t_i, t_j) = \frac{c\,\mathrm{Var}_i^+(t_i)}{\eta_j(t_j)\mu_j^+(t_j)\mu_i^+(t_i)} \longrightarrow 0,$$

provided $\eta_j(t_j)\, \mu_j^+(t_j)$ does not vanish at a faster rate than $\mathrm{Var}_i^+(t_i)$. We now show this is the case.

From Step 1 and the bound $V_j^*(t_j) \le V_j^*(0) \le R\hat{r}_j^\infty$, the FOC $V_j^*(t_j)\hat{\Lambda}_i(t_i) = c$ gives $\hat{\Lambda}_i(t_i) = \eta_i \mu_i^+ \ge c/(R\hat{r}_j^\infty) = \Omega(1/R)$, and similarly for j. In particular, $\eta_i, \eta_j = \Omega(1/R)$, i.e., they do not vanish faster than $1/R$.

Under the gap assumption, the convergence $\mathrm{Var}_i^+(t_i) \to 0$ is exponentially fast in t_i: writing the positive part of the prior as a sum of an atom at λ_{i1} with mass $p_{i,1}$ and a remainder supported on $(\lambda_{i1}+\varepsilon, \infty)$ for some $\varepsilon > 0$,

$$\mathrm{Var}_i^+(t_i) = O\big(e^{-\varepsilon\, t_i}\big) = O\big(R^{-\varepsilon/\lambda_{i1}}\big),$$

since t_i scales as $\lambda_{i1}^{-1}\log R$ near the abandonment point (by the asymptotic of the single-component FOC). This decay is strictly faster than $\eta_i = \Omega(1/R)$ once $\varepsilon/\lambda_{i1} > 1$ (and is at least faster than $\eta_i\eta_j = \Omega(1/R^2)$ for any $\varepsilon > 0$ when $\lambda_{i1} \leq 1$; an arbitrary lower bound on ε is recovered by the gap assumption). In summary, $A(t_i, t_j) \to 0$ at a strictly faster rate than $|B(t_i)| = \eta_i(\mu_i^+)^2 |V_{i,0}| = \Omega(\eta_i)$.

Therefore $\mathrm{Num}_i(t_i, t_j) \to B(t_i) < 0$ as $R \to \infty$, and the slope $(V_i^*/\hat{\Lambda}_i)'(t_i)$ is strictly negative for R sufficiently large.

Step 5: Conclusion. Writing $X_k := V_k^* \,\mathrm{Var}_k$ and $Y_k := \hat{\Lambda}_k \,\mathrm{Cov}_k$ for $k \in \{i, j\}$, the Hessian determinant from the proof of Theorem 2 satisfies the identity

$$2H = \big(X_j + Y_j\big)\,\mathrm{Num}_i + \big(X_i + Y_i\big)\,\mathrm{Num}_j,$$

which is checked directly: $(X_j + Y_j)(X_i - Y_i) + (X_i + Y_i)(X_j - Y_j) = 2X_iX_j - 2Y_iY_j = 2(V_i^*V_j^*\,\mathrm{Var}_i\,\mathrm{Var}_j - \hat{\Lambda}_i\hat{\Lambda}_j\,\mathrm{Cov}_i\,\mathrm{Cov}_j) = 2H$. Both bracketed factors $X_k + Y_k = V_k^*\,\mathrm{Var}_k + \hat{\Lambda}_k\,\mathrm{Cov}_k$ are non-negative. By Step 4, both Num_i and Num_j are strictly negative at any interior candidate once R is sufficiently large (the argument is symmetric in $i \leftrightarrow j$), so $H < 0$ at every such candidate. A local interior maximum requires $H \geq 0$, so no interior maximum exists for $R > \bar{R}$. Every optimum of (14) therefore lies on the boundary $\{T_1 = 0\} \cup \{T_2 = 0\}$, which by definition is a corner solution. □

Proof of Proposition 8. The argument proceeds via the asymptotic expansion of $I_i(R)$ around $R = \infty$.

Step 1: Asymptotic of T_i^*. The single-component first-order condition at the prior is

$$R\,\mathbb{E}\big[\Lambda_i\, e^{-(\Lambda_i+\delta)T_i^*}\big] = c\,\mathbb{E}\big[e^{-(\Lambda_i+\delta)T_i^*}\big].$$

For $T_i^* \to \infty$ the right-hand side is dominated by the atom at 0: $\mathbb{E}[e^{-(\Lambda_i+\delta)T_i^*}] = \pi_{i0}e^{-\delta T_i^*}(1 + o(1))$, the remainder being at most $(1-\pi_{i0})e^{-(\lambda_{i1}+\delta)T_i^*}/\pi_{i0}$. For the left-hand side, define

$$a_i(T) := \mathbb{E}\big[\Lambda_i\, e^{-(\Lambda_i - \lambda_{i1})T}\,\mathbf{1}\{\Lambda_i > 0\}\big],$$

so that $\mathbb{E}[\Lambda_i e^{-(\Lambda_i+\delta)T_i^*}] = e^{-(\lambda_{i1}+\delta)T_i^*}\, a_i(T_i^*)$. By dominated convergence, $a_i(T)$ is bounded above by $\mathbb{E}[\Lambda_i \mathbf{1}\{\Lambda_i > 0\}] < \infty$, and bounded below by

$$a_i(T) \geq \lambda_{i1}\,\mathbb{P}\big(\lambda_{i1} \leq \Lambda_i \leq \lambda_{i1} + \varepsilon\big)\, e^{-\varepsilon T}, \qquad \text{for any } \varepsilon > 0,$$

which is strictly positive for every $\varepsilon > 0$ since $\lambda_{i1} \in \mathrm{supp}(\Lambda_i)$. In particular, $\log a_i(T) = O(\varepsilon T)$ as $T \to \infty$, and one may let $\varepsilon \to 0$ to conclude that the FOC $R e^{-(\lambda_{i1}+\delta)T_i^*} a_i(T_i^*) = c\pi_{i0} e^{-\delta T_i^*}(1+o(1))$ yields

$$e^{-\lambda_{i1} T_i^*} = \frac{c\,\pi_{i0}}{R\,a_i(T_i^*)}\big(1+o(1)\big), \qquad T_i^* = \frac{\log R}{\lambda_{i1}}\big(1+o(1)\big). \tag{EC.18}$$

In particular $e^{-\delta T_i^*} = R^{-\delta/\lambda_{i1}}\big(\frac{c\,\pi_{i0}}{a_i(T_i^*)}\big)^{\delta/\lambda_{i1}}(1+o(1))$.

Step 2: Asymptotics of $\hat{r}_i$ and $\hat{c}_i$. From Proposition 4,

$$1 - \hat{r}_i = \mathbb{E}\left[\frac{\delta}{\Lambda_i + \delta}\right] + \mathbb{E}\left[\frac{\Lambda_i}{\Lambda_i + \delta}\, e^{-(\Lambda_i+\delta)T_i^*}\right] = \delta\,\mathbb{E}\left[\frac{1}{\Lambda_i + \delta}\right] + O\big(R^{-1-\delta/\lambda_{i1}}\big),$$

because the second expectation, restricted to $\{\Lambda_i > 0\}$, is bounded by $e^{-(\lambda_{i1}+\delta)T_i^*} = O(R^{-1-\delta/\lambda_{i1}})$. For $\hat{c}_i$, separating the atom at 0 from the positive part,

$$\hat{c}_i = c\,\mathbb{E}\left[\frac{1}{\Lambda_i + \delta}\right] - \frac{c\,\pi_{i0}}{\delta}\, e^{-\delta T_i^*}\big(1+o(1)\big) + O\big(R^{-1-\delta/\lambda_{i1}}\big),$$

since the $\{\Lambda_i > 0\}$ contribution to $\mathbb{E}[e^{-(\Lambda_i+\delta)T_i^*}/(\Lambda_i+\delta)]$ is at most $e^{-(\lambda_{i1}+\delta)T_i^*}/(\lambda_{i1}+\delta)$. Substituting $e^{-\delta T_i^*}$ from (EC.18),

$$\hat{c}_i = c\,\mathbb{E}\left[\frac{1}{\Lambda_i + \delta}\right] - \frac{c\,\pi_{i0}}{\delta}\left(\frac{c\,\pi_{i0}}{a_i(T_i^*)}\right)^{\delta/\lambda_{i1}} R^{-\delta/\lambda_{i1}} + o\big(R^{-\delta/\lambda_{i1}}\big). \tag{EC.19}$$

Step 3: Expansion of I_i. Combining Step 2 and the geometric expansion $1/(a - b\epsilon) = (1/a)(1 + b\epsilon/a + o(\epsilon))$ with $\epsilon = R^{-\delta/\lambda_{i1}} \to 0$,

$$I_i(R) = \frac{\delta}{c} + \frac{\pi_{i0}}{c\,\mathbb{E}\big[\frac{1}{\Lambda_i+\delta}\big]}\left(\frac{c\,\pi_{i0}}{a_i(T_i^*)}\right)^{\delta/\lambda_{i1}} R^{-\delta/\lambda_{i1}} + o\big(R^{-\delta/\lambda_{i1}}\big). \tag{EC.20}$$

Step 4: Tier-1 comparison. Under symmetric costs the leading terms δ/c in (EC.20) are identical for $i = 1, 2$. The subleading correction in (EC.20) is strictly positive, decays at rate $R^{-\delta/\lambda_{i1}}$ up to the factor $\big(c\pi_{i0}/a_i(T_i^*)\big)^{\delta/\lambda_{i1}}$, which by the bounds on a_i from Step 1 satisfies

$$\big(c\pi_{i0}/\mathbb{E}[\Lambda_i]\big)^{\delta/\lambda_{i1}} \le \big(c\,\pi_{i0}/a_i(T_i^*)\big)^{\delta/\lambda_{i1}} \le \big(c\pi_{i0}\big)^{\delta/\lambda_{i1}}\big(\lambda_{i1}\mathbb{P}(\lambda_{i1} \le \Lambda_i \le \lambda_{i1}+\varepsilon)\big)^{-\delta/\lambda_{i1}} e^{\varepsilon\delta T_i^*/\lambda_{i1}}.$$

For any $\varepsilon > 0$, plugging $T_i^* \le \frac{\log R}{\lambda_{i1}}(1+o(1))$ yields an upper bound of order $R^{\varepsilon\delta/\lambda_{i1}^2}$, which is slowly varying. Hence the correction $I_i(R) - \delta/c$ has dominant order $R^{-\delta/\lambda_{i1}+\varepsilon'}$ from above and $R^{-\delta/\lambda_{i1}}$ from below, for any $\varepsilon' > 0$ chosen in advance.

If $\lambda_{i1} > \lambda_{j1}$, choose $\varepsilon' > 0$ small enough that $\delta/\lambda_{i1} + \varepsilon' < \delta/\lambda_{j1}$, which is possible since $\delta/\lambda_{i1} < \delta/\lambda_{j1}$ strictly. Then for R sufficiently large, $I_i(R) - \delta/c \ge C_i R^{-\delta/\lambda_{i1}}$ and $I_j(R) - \delta/c \le C_j R^{-\delta/\lambda_{i1}-\varepsilon'/2}$ for some positive constants C_i, C_j, so $I_i > I_j$. This proves Part 1.

Step 5: Tier-2 comparison under an atom at λ_{i1}. If both distributions place positive mass π_{*1} at $\lambda_{i1}=\lambda_{j1}=\lambda_1$, then by dominated convergence $a_i(T)\to\pi_{i1}\lambda_{i1}$ as $T\to\infty$, and the expansion (EC.20) becomes the explicit formula

$$I_i(R)=\frac{\delta}{c}+\frac{\pi_{i0}}{c\,\mathbb{E}_{\Lambda_i}\left[\frac{1}{\Lambda_i+\delta}\right]}\left(\frac{c\,\pi_{i0}}{\pi_{i1}\,\lambda_1}\right)^{\delta/\lambda_1}R^{-\delta/\lambda_1}+o\big(R^{-\delta/\lambda_1}\big).$$

The same expansion applies to I_j with $i\leftrightarrow j$. Both expansions share the same R-power, so the comparison $I_i>I_j$ for large R reduces to comparison of the coefficients, which up to the common factor $1/c$ are exactly (27). This proves Part 2. $\square$

Proof of Proposition 9. The proof builds on the asymptotic expansion of $V_i^*(\rho)$ as $\rho\to\infty$, derived next, and on the index expansion (EC.20) of Proposition 8.

Step 1: Asymptotic expansion of $V_i^*(\rho)$. By the single-component formula (3),

$$V_i^*(\rho)=\mathbb{E}\left[\frac{\Lambda_i\rho-c}{\Lambda_i+\delta}\big(1-e^{-(\Lambda_i+\delta)T_i^*(\rho)}\big)\right]=V_i^{(\infty)}(\rho)-\mathbb{E}\left[\frac{\Lambda_i\rho-c}{\Lambda_i+\delta}\,e^{-(\Lambda_i+\delta)T_i^*(\rho)}\right],$$

where $V_i^{(\infty)}(\rho):=\rho\,\hat r_i^\infty-c\,\mathbb{E}[1/(\Lambda_i+\delta)]$ denotes the limit as $T_i^*\to\infty$ (and is itself the prior value of a policy that never abandons). By the same Step 1 argument as in the proof of Proposition 8, $T_i^*(\rho)=\lambda_{i1}^{-1}\log\rho\,(1+o(1))$ and the dominant contributions to the second expectation come from $\Lambda_i=0$ (for the $-c$ piece) and from the neighborhood of $\Lambda_i=\lambda_{i1}$ (for the $\Lambda_i\rho$ piece). Collecting the leading terms,

$$V_i^*(\rho)=\rho\,\hat r_i^\infty-c\,\mathbb{E}\left[\frac{1}{\Lambda_i+\delta}\right]+\gamma_i\big(\rho\big)\,\rho^{-\delta/\lambda_{i1}}+o\big(\rho^{-\delta/\lambda_{i1}}\big),\tag{EC.21}$$

where the coefficient $\gamma_i(\rho):=\frac{c\pi_{i0}\lambda_{i1}}{\delta(\lambda_{i1}+\delta)}\big(c\pi_{i0}/a_i(T_i^*(\rho))\big)^{\delta/\lambda_{i1}}$ is strictly positive and slowly varying in ρ (with a_i defined as in the proof of Proposition 8). In the discrete case with an atom $\pi_{i1}>0$ at λ_{i1}, $a_i(T)\to\pi_{i1}\lambda_{i1}$ and γ_i converges to the explicit constant $\frac{c\pi_{i0}\lambda_{i1}}{\delta(\lambda_{i1}+\delta)}\big(c\pi_{i0}/(\pi_{i1}\lambda_{i1})\big)^{\delta/\lambda_{i1}}$.

Step 2: Asymptotic expansion of V_i^{**}. By definition, $V_i^{**}=V_i^*\big(V_j^*(R)\big)$. Applying (EC.21) twice, with the outer evaluation at $\rho=V_j^*(R)=R\,\hat r_j^\infty-c\,\mathbb{E}[1/(\Lambda_j+\delta)]+\gamma_j(R)R^{-\delta/\lambda_{j1}}+\cdots$, gives

$$\begin{aligned}V_i^{**}&=V_j^*(R)\,\hat r_i^\infty-c\,\mathbb{E}\left[\frac{1}{\Lambda_i+\delta}\right]+\gamma_i(V_j^*(R))\big[V_j^*(R)\big]^{-\delta/\lambda_{i1}}+\cdots\\&=R\,\hat r_i^\infty\hat r_j^\infty-c\,\hat r_i^\infty\mathbb{E}\left[\frac{1}{\Lambda_j+\delta}\right]-c\,\mathbb{E}\left[\frac{1}{\Lambda_i+\delta}\right]\\&\quad+\hat r_i^\infty\,\gamma_j(R)\,R^{-\delta/\lambda_{j1}}+\gamma_i\big(V_j^*(R)\big)\big(\hat r_j^\infty\big)^{-\delta/\lambda_{i1}}R^{-\delta/\lambda_{i1}}+o\big(R^{-\min(\delta/\lambda_{i1},\delta/\lambda_{j1})}\big).\end{aligned}$$

The $O(R)$ and $O(1)$ terms in V_i^{**} are *symmetric in* (i,j), which one verifies directly using the identity $1-\hat r_k^\infty=\delta\,\mathbb{E}[1/(\Lambda_k+\delta)]$:

$$\big(V_i^{**}-V_j^{**}\big)_{R\text{-term}}=0,\qquad\big(V_i^{**}-V_j^{**}\big)_{O(1)\text{-term}}=c\,\delta\big[\mathbb{E}\tfrac{1}{\Lambda_j+\delta}\mathbb{E}\tfrac{1}{\Lambda_i+\delta}-\mathbb{E}\tfrac{1}{\Lambda_i+\delta}\mathbb{E}\tfrac{1}{\Lambda_j+\delta}\big]=0.\tag{EC.22}$$

Hence the ordering of V_1^{**} versus V_2^{**} at large R is governed entirely by the four subleading terms above.

Step 3: Comparison under $\lambda_{i1} \neq \lambda_{j1}$. Without loss of generality assume $\lambda_{i1} > \lambda_{j1}$, so that $\delta/\lambda_{i1} < \delta/\lambda_{j1}$ and the $R^{-\delta/\lambda_{i1}}$ contributions dominate the $R^{-\delta/\lambda_{j1}}$ contributions (the latter decay strictly faster). Collecting the $R^{-\delta/\lambda_{i1}}$ coefficients from Step 2 (the $\gamma_i\big(V_j^*(R)\big)$ term in V_i^{**} and the $\hat{r}_j^{\infty}\,\gamma_i(R)$ term in V_j^{**}, the latter inherited from the $R^{-\delta/\lambda_{i1}}$ correction in $V_i^*(R)$ that is then multiplied through by the leading $\hat{r}_j^{\infty}$ in V_j^*),

$$\big(V_i^{**} - V_j^{**}\big) = \Big[\gamma_i\big(V_j^*(R)\big)\big(\hat{r}_j^{\infty}\big)^{-\delta/\lambda_{i1}} - \hat{r}_j^{\infty}\,\gamma_i(R)\Big]R^{-\delta/\lambda_{i1}} + o\big(R^{-\delta/\lambda_{i1}}\big). \tag{EC.23}$$

Because $\gamma_i(\rho) = \frac{c\pi_{i0}\lambda_{i1}}{\delta(\lambda_{i1}+\delta)}\big(c\pi_{i0}/a_i(T_i^*(\rho))\big)^{\delta/\lambda_{i1}}$ is slowly varying in ρ (a function of $T_i^*(\rho) \sim \lambda_{i1}^{-1}\log\rho$), the ratio $\gamma_i(V_j^*(R))/\gamma_i(R) \to 1$ as $R \to \infty$. Hence the bracket in (EC.23) is asymptotic to $\gamma_i(R)\big[\big(\hat{r}_j^{\infty}\big)^{-\delta/\lambda_{i1}} - \hat{r}_j^{\infty}\big]$, which is strictly positive because $\gamma_i(R) > 0$ for all R and the map $x \mapsto x^{-\delta/\lambda_{i1}} - x$ is strictly positive on $(0,1)$ (since $x^{-\delta/\lambda_{i1}} > 1 > x$), where $\hat{r}_j^{\infty} \in (0,1)$ by assumption. Hence $V_i^{**} > V_j^{**}$ for all R sufficiently large, and the optimal policy of Proposition 7 picks component i.

By Part 1 of Proposition 8, the index also picks component i for all R sufficiently large. Therefore the two selections coincide, and the modified index policy, which by construction attains $\max_k V_k^{**}$ once it picks the correct component, satisfies $\tilde{V}(I) = V^{**}$. This proves Part 2.

Step 4: Asymptotic optimality (Part 1). For either component $k \in \{1,2\}$, Step 2 gives $V_k^{**}/R \to \hat{r}_1^{\infty}\hat{r}_2^{\infty}$ as $R \to \infty$ (the leading R-coefficient is $\hat{r}_1^{\infty}\hat{r}_2^{\infty}$ irrespective of which component is sequenced first, because of the cancellation (EC.22)). Hence $V^{**}/R = \max_k V_k^{**}/R \to \hat{r}_1^{\infty}\hat{r}_2^{\infty}$. The modified index policy attains $\tilde{V}(I) = V_i^{**}$ for i the index-selected component, so $\tilde{V}(I)/R \to \hat{r}_1^{\infty}\hat{r}_2^{\infty}$ as well, regardless of which side the index picks; the limit $\tilde{V}(I)/V^{**} \to 1$ follows.

For the original index policy, $V(I)$ equals $V(T_i^*, 0)$ for i the index-selected component. By the expression for $V(T_i, 0)$ in (13) (or the calculation in the proof of Proposition 5), $V(T_i^*, 0) = \hat{r}_i\,V_j^*(R) - \hat{c}_i$, where $\hat{r}_i, \hat{c}_i$ are the index quantities of Proposition 4. From the expansions $\hat{r}_i \to \hat{r}_i^{\infty}$ and $\hat{c}_i/R \to 0$ as $R \to \infty$ (both established in the proof of Proposition 8), and $V_j^*(R)/R \to \hat{r}_j^{\infty}$ from Step 1, $V(T_i^*, 0)/R \to \hat{r}_i^{\infty}\hat{r}_j^{\infty}$. The same limit holds for $V(0, T_j^*)$ by symmetry, so $V(I)/R \to \hat{r}_i^{\infty}\hat{r}_j^{\infty}$ and $V(I)/V^{**} \to 1$. $\square$

EC.3. Additional Results.

EC.3.1. The no-switching result for the $n > 2$ components

Theorem 3 (n-component no-switching) . *For each* $i \in [n] := \{1, \dots, n\}$*, suppose the ratio* $V_i^*(t_i)/\hat{\Lambda}_i(t_i)$ *is* strictly decreasing *in* t_i *on* $[0, T_i^*)$*. Then the n-component value function* $V_n(T_1, \dots, T_n)$ *admits no interior local maximum on* $\prod_i [0, T_i^*]$*; consequently every optimum is a corner.*

Proof. We argue by induction on n. The base case $n = 2$ is Theorem 2, whose proof gives the strict conclusion $H < 0$ at every interior critical point under the strict slope condition.

For the induction step, fix $n \geq 3$ and assume the theorem for $2 \leq k < n$. Let $\bar{T} = (\bar{T}_1, \dots, \bar{T}_n)$ be an interior critical point of V_n. We translate posteriors so that the posteriors at $\bar{T}_k$ become the new priors and the critical point sits at the origin; this is the same device used in the proof of Theorem 2 and changes no statement of the slope condition. Hereafter all expectations $\hat{\Lambda}_k$, variances $\mathrm{Var}(\Lambda_k)$, and continuation values refer to these translated priors.

Step 1: Hessian entries. Differentiating V_n twice at the origin (the calculation is the direct n-component analogue of the one in the proof of Theorem 2), we obtain

$$H_n[i,i] = -\rho_i \ \mathrm{Var}(\Lambda_i), \qquad H_n[i,j] = \hat{\Lambda}_j \frac{\partial \rho_j}{\partial T_i} = \hat{\Lambda}_i \frac{\partial \rho_i}{\partial T_j} \quad (i \neq j), \tag{EC.24}$$

where $\rho_k := V^*_{[n]\setminus\{k\}}$ is the optimal value of the $(n-1)$-component sub-problem on $[n] \setminus \{k\}$, and the second equality in the off-diagonal is Schwarz's identity (verified directly by computing the two sides from the analogue of (EC.2)). Note that ρ_k does not depend on T_k: the sub-problem $[n] \setminus \{k\}$ does not include component k.

Step 2: First-order conditions. Differentiating V_n once at the origin gives, by the same calculation that yields (EC.5) in the proof of Theorem 2, $\partial V_n/\partial T_i\big|_{\bar{T}} = \hat{\Lambda}_i \, \rho_i - c$. The interior FOC $\partial V_n/\partial T_i = 0$ therefore gives

$$\rho_i = \frac{c}{\hat{\Lambda}_i}, \qquad i = 1, \dots, n. \tag{EC.25}$$

Step 3: A signed bound on the off-diagonal entry. By the inductive hypothesis applied to the $(n-1)$-component sub-problem on $[n] \setminus \{i\}$, that sub-problem admits no interior optimum, so its optimal value is attained on a corner: there exists $\sigma(i) \in [n] \setminus \{i\}$ such that

$$\rho_i = V^*_{\sigma(i)}\big(V^*_{[n]\setminus\{i,\sigma(i)\}}\big), \tag{EC.26}$$

where $V^*_{\sigma(i)}(\rho)$ denotes the single-component optimum on component $\sigma(i)$ with continuation reward ρ. Combining (EC.26) with the second form of $H_n[i, \sigma(i)]$ in (EC.24) and the fact that ρ_i depends on $T_{\sigma(i)}$ only through the single-component value $V^*_{\sigma(i)}$ at posterior age $T_{\sigma(i)}$, we have

$$H_n[i, \sigma(i)] = \hat{\Lambda}_i \frac{\partial \rho_i}{\partial T_{\sigma(i)}} = \hat{\Lambda}_i \frac{dV^*_{\sigma(i)}}{dT_{\sigma(i)}}, \tag{EC.27}$$

the inner derivative being the single-component slope of $V^*_{\sigma(i)}$ with respect to its own posterior age.

We now apply the strict slope condition to component $\sigma(i)$. By hypothesis, the ratio $V^*_{\sigma(i)}(t)/\hat{\Lambda}_{\sigma(i)}(t)$ is strictly decreasing in t; equivalently, by the quotient rule and the identity $d\hat{\Lambda}_k(t)/dt = -\operatorname{Var}(\Lambda_k(t))$ (Lemma 2),

$$\hat{\Lambda}_{\sigma(i)} \frac{dV^*_{\sigma(i)}}{dT_{\sigma(i)}} < -V^*_{\sigma(i)}\big(V^*_{[n]\setminus\{i,\sigma(i)\}}\big) \operatorname{Var}(\Lambda_{\sigma(i)}) = -\rho_i \operatorname{Var}(\Lambda_{\sigma(i)}), \tag{EC.28}$$

where the last equality uses (EC.26). Both sides of (EC.28) are negative ($dV^*_{\sigma(i)}/dT_{\sigma(i)} < 0$ on $[0, T^*_{\sigma(i)})$, and $\rho_i, \operatorname{Var}(\Lambda_{\sigma(i)}) > 0$), which fixes the direction of the inequality unambiguously.

Multiplying (EC.28) through by $\hat{\Lambda}_i/\hat{\Lambda}_{\sigma(i)} > 0$ preserves the direction, and substituting into (EC.27) yields

$$H_n[i,\sigma(i)] = \hat{\Lambda}_i \frac{dV^*_{\sigma(i)}}{dT_{\sigma(i)}} < -\frac{\hat{\Lambda}_i\,\rho_i \operatorname{Var}(\Lambda_{\sigma(i)})}{\hat{\Lambda}_{\sigma(i)}} \overset{(EC.25)}{=} -\frac{c \operatorname{Var}(\Lambda_{\sigma(i)})}{\hat{\Lambda}_{\sigma(i)}} \overset{(EC.25)}{=} -\rho_{\sigma(i)} \operatorname{Var}(\Lambda_{\sigma(i)}). \tag{EC.29}$$

In particular $H_n[i,\sigma(i)] < 0$, and the strict negativity has the sharp magnitude $|H_n[i,\sigma(i)]| > \rho_{\sigma(i)} \operatorname{Var}(\Lambda_{\sigma(i)})$.

Step 4: A pigeonhole on dispersion ratios. Define $r_k := \operatorname{Var}(\Lambda_k)/\hat{\Lambda}_k > 0$. We claim there exists $i \in [n]$ such that $r_{\sigma(i)} \geq r_i$. Indeed, suppose for contradiction that $r_{\sigma(i)} < r_i$ for every i. Form the orbit $i_0, i_1 = \sigma(i_0), i_2 = \sigma(i_1), \ldots$. By the contradiction hypothesis the sequence $r_{i_0}, r_{i_1}, r_{i_2}, \ldots$ is strictly decreasing. Since the orbit takes values in the finite set $[n]$, it eventually enters a cycle, on which the r_{i_k} values must repeat; but a repeated value in a strictly decreasing sequence is impossible.

Step 5: Conclusion. Pick i from Step 4 so that $r_{\sigma(i)} \geq r_i$, equivalently $\rho_{\sigma(i)} \operatorname{Var}(\Lambda_{\sigma(i)}) \geq \rho_i \operatorname{Var}(\Lambda_i)$ (using (EC.25), this rearranges to $r_{\sigma(i)} \geq r_i$). Let $j := \sigma(i)$ and take the vector $v = e_i - e_j \in \mathbb{R}^n$. Then

$$\begin{aligned} v^\top H_n v &= H_n[i,i] - 2H_n[i,j] + H_n[j,j] \\ &= -\rho_i \operatorname{Var}(\Lambda_i) - 2H_n[i,j] - \rho_j \operatorname{Var}(\Lambda_j). \end{aligned}$$

From (EC.29), $H_n[i,j] < -\rho_j \operatorname{Var}(\Lambda_j)$, so $-2H_n[i,j] > 2\rho_j \operatorname{Var}(\Lambda_j)$. Hence

$$v^\top H_n v > -\rho_i \operatorname{Var}(\Lambda_i) + 2\rho_j \operatorname{Var}(\Lambda_j) - \rho_j \operatorname{Var}(\Lambda_j) = \rho_j \operatorname{Var}(\Lambda_j) - \rho_i \operatorname{Var}(\Lambda_i) \geq 0,$$

where the final inequality is the pigeonhole choice $\rho_j \operatorname{Var}(\Lambda_j) \geq \rho_i \operatorname{Var}(\Lambda_i)$. Combined with the strict inequality from (EC.29), we conclude $v^\top H_n v > 0$. Hence H_n has at least one positive eigenvalue at $\bar{T}$, so $\bar{T}$ is not a local maximum. □

The result gives two immediate corollaries.

Corollary 4 (n-component, two-point priors) . *When each Λ_i has two-point support $\{0, \lambda_i\}$, the slope ratio $V_i^*(t_i)/\hat{\Lambda}_i(t_i)$ is strictly decreasing on $[0, T_i^*)$ (the calculation is the one in the proof of Proposition 2); hence V_n admits only corner solutions for every n.*

Corollary 5 (n-component moonshot under feasibility uncertainty) . *If each Λ_i satisfies $\mathbb{P}(\Lambda_i = 0) > 0$, $\inf\big(\mathrm{supp}(\Lambda_i) \cap (0,\infty)\big) > 0$, and $\mathbb{E}[\Lambda_i] < \infty$, then by the asymptotic argument of Proposition 7 the strict slope condition holds for every component once R is sufficiently large. Hence V_n admits only corner solutions for R large.*

Some remarks on why the new pigeonhole step appears in the proof. The map $\sigma : [n] \to [n]$ from Step 3 records, for each component i, the inductively-optimal first component of the sub-problem on $[n] \setminus \{i\}$. When σ has a 2-cycle ($\sigma(i) = j$ and $\sigma(j) = i$), one can show that the 2×2 principal submatrix on (i, j) has non-positive determinant directly, and that submatrix is indefinite. When σ has only longer cycles, no individual 2×2 submatrix is automatically indefinite, but Step 4 guarantees that *some* pair $(i, \sigma(i))$ satisfies $r_{\sigma(i)} \geq r_i$, and the strict slope condition is what converts this into a strict positive eigenvalue of the full Hessian via the vector $v = e_i - e_{\sigma(i)}$.

EC.3.2. Performance with more than two components

The model and the index policy in the main paper extend to projects with $n > 2$ components in the obvious way. Let $S \subseteq \{1, \ldots, n\}$ denote the set of components not yet attempted. The dynamic-programming recursion

$$V^*(S) \;=\; \max\Big\{0,\ \max_{i \in S} V_i^*\big(S \setminus \{i\}\big)\Big\}, \qquad V^*(\emptyset) = R, \tag{EC.30}$$

where $V_i^*(\cdot)$ denotes the single-component optimal value of working on component i with continuation reward $V^*(S \setminus \{i\})$, yields the optimal value $V^*(\{1, \ldots, n\})$ by enumeration over 2^n subsets. The natural multi-component extension of the modified index policy of Corollary 3 is to rank the components in decreasing order of I_i (computed against the reference reward R as before) and to evaluate the value of working on them in that fixed order via the same nested V_i^* recursion.

We compare the two on Monte Carlo samples with 30,000 samples per (n, R) cell at $n \leq 4$, scaling down at large n where the subset DP is more expensive. To ensure that the optimal value is positive with non-negligible probability even at large n (the project pays R only if every component succeeds, so $\prod_i p_i$ shrinks geometrically in n), we sample $p_i \sim \mathcal{U}(0.5, 0.99)$, $\log \lambda_i \sim \mathcal{U}(\log 0.5, \log 5)$, $\log c \sim \mathcal{U}(\log 0.1, \log 2)$, $\delta \sim \mathcal{U}(0.05, 0.5)$. Figure EC.1 reports the mean relative and mean absolute error of the modified index policy as a function of n for five values of R.

Three features stand out.

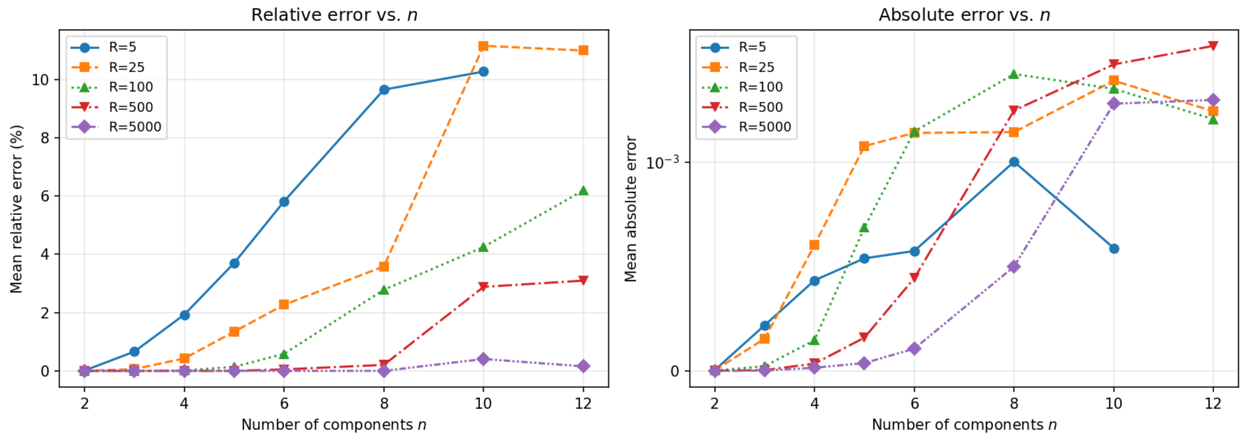

Figure EC.1 **Mean error of the modified-index policy as a function of the number of components n, for five values of the reward R. Left: relative error $(V^* - \tilde{V}(I))/V^*$. Right: absolute error $V^* - \tilde{V}(I)$ on a symlog axis. Each curve terminates at the largest n for which at least 10 parameter draws had $V^* > 0$.**

First, *relative error grows smoothly with n at fixed R, and shrinks rapidly in R at fixed n.* At R=500 the mean relative error stays below 3.1% through n=12, and at R=5000 it never exceeds 0.2% on the same range. The picture is monotone in both directions, and the curves for the smaller rewards rise above the others in the order one would expect: small reward and many components make the ordering decision matter most.

Second, *the moonshot pattern of Proposition 6 survives the extension to $n > 2$.* The R=5000 curve in the left panel hugs zero across the entire range, and even at n=12 the mean relative error is only 0.16%. As $R \to \infty$ each component's overshoot loss vanishes individually, and the index correctly orders the components regardless of how many of them there are.

Third, the right panel of Figure EC.1 shows that *absolute error saturates with n rather than diverging*. Across all five R values the mean absolute loss caps out near 10^{-3} to 10^{-2} in raw value units even at n=12. The large relative-error numbers at small R are therefore an artifact of the optimal value itself being small in those cells: the dollar amount the modified index forgoes is uniformly modest, even when, expressed as a fraction of V^*, it can run into double digits. The mechanism is the same $|O_1 - O_2| - |\Phi|$ bound as in (25), summed across the n candidate first-components: every parameter point where the index puts a "should-not-be-started-first" component at the top of the ordering incurs at most $|O_1 - O_2|$ of error on the relevant pair, and the overshoot losses O_i scale with single-component primitives, not with n.

Taken together with the closed-form bound for the two-component case, these experiments suggest that the modified index policy is a practically useful heuristic well beyond the two-component setting in which we have analytic guarantees: it is essentially exact in the moonshot regime for any n we tested, degrades gracefully in relative terms as n grows at moderate R, and incurs absolute losses that remain bounded in n on the parameter range we explored.

EC.3.3. An $O(n \cdot 2^n)$ algorithm for exact optimum in the $n > 2$ model.

Algorithm 1 Subset DP for the n-component optimum.

```
1: Initialize V[∅] ← R.
2: for k = 1 to n do
3:     for all subsets S ⊆ {1,...,n} with |S| = k do
4:         V[S] ← max{0, max_{i∈S} V_i^*(V[S \ {i}])}
5: return V[{1,...,n}].
```

The algorithm computes $V^*(\mathbf{S})$ for every $\mathbf{S} \subseteq \{1,\ldots,n\}$ by processing subsets in order of increasing cardinality, so that the values $V[\mathbf{S} \setminus \{i\}]$ needed in line 4 are already available when $V[\mathbf{S}]$ is computed. Since the single-component value function $V_i^*(\cdot)$ has the closed-form expression given by inserting (7) into (3), each evaluation in line 4 is $O(1)$, and computing $V[\mathbf{S}]$ for a subset of cardinality k costs $O(k)$. The total work is therefore

$$\sum_{k=0}^{n} k\binom{n}{k} = n \cdot 2^{n-1} = O\big(n \cdot 2^n\big),$$

strictly faster than the $O(n \cdot n!)$ cost of enumerating permutations of $\{1,\ldots,n\}$.

EC.3.4. Why the polynomial-decay case is problematic in Proposition 7.

For a prior whose positive part has continuous density $f^+(\lambda)$ with $f^+(\lambda_{i1}^+) = f_0 > 0$, the Laplace-method expansion at $t_i \to \infty$ gives

$$\mathbb{E}\big[\Lambda_i \,\big|\, \Lambda_i > 0, \tau_i > t_i\big] = \lambda_{i1} + \frac{1}{t_i} + o(1/t_i),$$

$$\mathrm{Var}\big[\Lambda_i \,\big|\, \Lambda_i > 0, \tau_i > t_i\big] = \frac{1}{t_i^2} + o(1/t_i^2).$$

Substituting $t_i \sim \lambda_{i1}^{-1} \log R$ (which is the leading asymptotic of the interior FOC location in Step 1) gives $\mathrm{Var}_i^+ \sim 1/(\log R)^2$. Combined with $\eta_i = O(1/R)$ (up to slowly varying factors), the leading term A in the decomposition (EC.17), $A = c\,\mathrm{Var}_i^+/(\eta_j \mu_j^+ \mu_i^+)$, scales as $O\big(R/(\log R)^2\big)$ and dominates $|B| = O(1/R)$. Hence $\mathrm{Num}_i > 0$ at the interior FOC, the slope of $V_i^*/\hat{\Lambda}_i$ is positive there, and the proof of Theorem 2 no longer rules out an interior maximum. A direct check of the Hessian determinant $H = V_1^* V_2^* \mathrm{Var}_1 \mathrm{Var}_2 - \hat{\Lambda}_1 \hat{\Lambda}_2 \mathrm{Cov}_1 \mathrm{Cov}_2$ under these scalings gives $H = O(c^2) > 0$ for R large, confirming that an interior local maximum can in fact exist.

If there is an atom at λ_{i1} then $\eta\mathrm{Var}_i^+ \to \pi_{i1}\lambda_{i1}^2 - \pi_{i1}^2\lambda_{i1}^2$ on the positive part (which together with the atom-driven exponential decay of the off-atom mass gives $\mathrm{Var}_i^+ = O(e^{-\varepsilon t_i})$). The alternate assumption of a spectral gap $(\inf\big(\mathrm{supp}(\Lambda_i) \cap (\lambda_{i1}, \infty)\big) > \lambda_{i1})$ of size $\eta > 0$ forces $\mathrm{Var}_i^+ \le \mathrm{Var}^+(\lambda_{i1}, \lambda_{i1} + \eta) + e^{-\eta t_i} \cdot (\text{bounded}) = O(e^{-\eta t_i})$. If either $\mathbb{P}(\Lambda_i = \lambda_{i1}) > 0$ or $\inf\big(\mathrm{supp}(\Lambda_i) \cap (\lambda_{i1}, \infty)\big) > \lambda_{i1}$, $\mathrm{Var}_i^+(t_i)$ decays at least exponentially in t_i, so the Step 4 dominance argument goes through unchanged.

EC.3.5. Full table of results on the performance of modified index policy on the robust subset

		Relative error on robust subset		
R	#robust	mean	99th %	99.9th %
0.1	226	0.00%	0.0%	0.0%
0.2	587	0.01%	0.0%	1.6%
0.5	1,532	0.02%	0.0%	4.8%
1	2,654	0.02%	0.0%	5.3%
2	4,289	0.02%	0.0%	6.0%
5	7,220	0.03%	0.0%	8.7%
10	9,943	0.03%	0.0%	5.5%
25	13,536	0.03%	0.0%	6.2%
50	16,683	0.01%	0.0%	3.5%
100	19,484	0.03%	0.0%	7.1%
500	24,936	0.02%	0.0%	4.8%
5,000	28,817	0.01%	0.0%	0.6%

Table EC.1 **Same Monte Carlo as Table 1, on the robust subset $V_{\mathrm{opt}} > 0.05\,c$ (so that relative error is a meaningful normalization).**