Budgeting Discretion: Theory and Evidence on Street-Level Decision-Making

Gaurab Pokharel
Institute for Advanced Computing
Virginia Tech
Alexandria, VA 22305
gaurab@vt.edu

Sanmay Das
Institute for Advanced Computing
Virginia Tech
Alexandria, VA 22305
sanmay@vt.edu

Patrick J. Fowler
Bursky School of Public Health
Washington University in Saint Louis
St. Louis, MO 63130
pjfowler@wustl.edu

Abstract

Street-level bureaucrats, such as caseworkers, triage nurses, and school administrators, operate under a fundamental tension. The official policy on who to prioritize for services may conflict with bureaucrats' estimation of who would benefit most from those services. The social theory of street-level bureaucracy argues that frontline agents use discretionary authority to bridge this gap, overriding policy recommendations when professional judgment identifies cases where doing so would meaningfully improve outcomes. As algorithmic recommendations become embedded in allocation settings, understanding the structure of this discretionary behavior is fundamental to designing tools that account for, rather than suppress, this value of professional judgment. We formalize discretion as a finite-horizon dynamic allocation problem. An agent implements a default policy but holds a limited budget of K overrides across T periods. Each period, the agent observes whether deviating from the default would yield a welfare gain and must decide whether to spend a scarce override now or conserve it for future opportunities. We show that the optimal policy is a state-dependent threshold rule and prove a behavioral invariance: for location-scale families of improvement distributions, the *rate* at which an optimal agent exercises discretion is independent of the amount of welfare gains and depends only on the distribution's shape. Fat-tailed gains induce patience (conserving overrides for rare, high-stakes cases); thin-tailed gains induce routine, front-loaded spending. We also compare the conceptual implications of our model with operational data from a homelessness service system. In the model, a capacity opening makes an override cheaper to spend now, which lowers the threshold for using one. In the operational data, override rates indeed shift with short-run capacity openings in scarce housing, rising when recent exits free slots, and falling when interventions are congested. At a descriptive level, we also find that the temporal pattern of spending is front-loaded: override rates peak on Mondays and early in the fiscal year. Together, the model and the data point to a systematic structure in how discretion is used, with direct implications for the design of human-AI allocation systems that are both more efficient and equitable.

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

1 Introduction

Frontline public-service workers often operate under two directives that can pull in opposite directions. Policies often mandate prioritizing the most vulnerable, whether that means the household with the highest acuity score, the patient with the most urgent need, or the child at greatest risk [16, 43]. But the outcomes on which these same workers are evaluated often favor whoever would benefit most from the intervention, whether that is the household that will stay stably housed [23], the patient whose prognosis improves the most from treatment [30], or the family most responsive to services. The most vulnerable person and the person who would benefit most are not always the same person. The theory of street-level bureaucracy has long recognized that caseworkers use professional discretion to bridge this gap, drawing on contextual knowledge and local judgment that no centralized rule can fully encode [1, 8, 9, 25].

When caseworkers act on this judgment, they deviate from whatever standing rule governs the allocation – a scoring instrument, a formulaic heuristic, or increasingly, an algorithmic recommendation [33, 37, 40, 41]. But discretion is not free. Prior work on discretionary behavior, whether in the street-level bureaucracy tradition [6, 25] or in the growing literature on expert responses to algorithmic advice [1, 11, 14, 20], has largely treated each override as a per-case choice.

In practice, override capacity is finite, constrained by downstream resources (e.g., housing slots, ICU beds), managerial oversight, and operational bandwidth [4, 13, 25], and using an override today reduces the ability to do so tomorrow. The classic theory of street-level bureaucracy emphasizes precisely this point: discretion is shaped by caseload pressure and institutional scarcity [4, 6, 25]. Yet, formal models of how finite discretion should be rationed over time are largely absent [32, 33]. Understanding the temporal structure of override behavior is a prerequisite for designing algorithmic tools that support professional judgment rather than suppress it.

We address this gap by formalizing the problem of *Budgeted Discretion*. We ground the model in homelessness services, where caseworkers assign households to one of two interventions: Emergency Shelter (ES), a lower-intensity default, or Transitional Housing (TH), a scarce but stabilizing alternative [23, 33]. An agent has a budget of K overrides over a finite horizon of T periods. In each period t, the default algorithmic recommendation assigns an action (e.g., assign ES rather than TH), and the agent observes whether overriding that recommendation would improve welfare by an incremental amount Δ_t (the *gain from discretion*). The gain Δ_t can either be positive (the agent can strictly improve the outcome by overriding) or effectively zero (the default recommendation already matches judgment).[1] If the agent overrides, they spend one unit of budget and realize the improvement; if not, they conserve budget for future opportunities. In our domain of homelessness services, the caseworker can override the default to move a household from ES to TH. This framing connects sociological accounts of discretion under scarcity to a tractable finite-horizon control problem closely related to online selection with capacity constraints [2, 18, 22].

Our main theoretical contribution is a *behavioral invariance*. We prove that, for location-scale families of improvement distributions [5], the rate at which an optimal agent exercises discretion at every state of the dynamic program is invariant to affine rescaling of payoffs and depends only on the distribution's shape. Intuitively, changing units (e.g., dollars versus thousands of dollars) rescales both the gains and the opportunity cost of spending by the same factor, so the probability that the agent overrides in any given state is unchanged. Once scale drops out, the key remaining driver of behavior is tail heaviness. With a fat-tailed gain distribution, most override opportunities yield modest improvements but rare cases produce very large gains, as when a housing override averts an imminent safety crisis. In this setting, optimal agents are patient, conserving overrides for outlier opportunities because the option value of waiting is high. With a thin-tailed distribution, improvements are relatively uniform, and agents override earlier and more routinely because waiting is unlikely to reveal a dramatically better case.

We compare the qualitative implications of our model with administrative data from the St. Louis Homeless Management Information System (HMIS), where caseworkers can deviate from an (algorithmically recovered) baseline assignment rule. We find two major effects in the data. First, when recent exits free slots in an intervention, caseworkers shift their overrides toward that intervention. In the model, this corresponds to a lower threshold for spending an override: when slots are open, each override costs less and draws less scrutiny.Second, caseworkers exercise more discretion early in the

[1] Δ_t is induced by a baseline welfare draw and an independent "improvement" draw (more on this in Section 2).

calendar, both within the week and within the fiscal year. The model predicts such front-loading under thin-tailed gains, but in these data staffing and intake schedules could produce it too, so these results should be taken as descriptive. Together, these findings identify systematic structure in observed override behavior, with direct implications for the design of human-AI allocation systems that support effective frontline decision-making.

Relationship to prior work: Our model of budgeted discretion sits at the intersection of several literatures on sequential decision-making under scarcity. Here, we briefly trace the connection from budgeted allocation to threshold-based stopping, then to invariance arguments, and finally to the applied discretion literature. A detailed discussion appears in Appendix B.

Structurally, our problem is online allocation with limited budget: each period reveals a stochastic opportunity (the incremental gain from overriding a default), and the agent decides whether to spend one unit of a finite budget or save it. This parallels dynamic and stochastic knapsack models [22, 10, 2] and prophet inequality benchmarks [17, 18, 21]. In contrast to this value- and regret-focused literature, we characterize *the rate* at which discretionary capacity is consumed. This lens connects to optimal stopping and search, where reservation values balance current gains against the option value of waiting [26, 44], and to multi-choice secretary problems with thresholds indexed by remaining selections [12, 35, 42]. In our work, the system always implements a baseline policy, and the scarce action is an override whose payoff is the incremental gain Δ_t relative to that baseline. Our invariance result connects to a tradition of parameter-free sequential decision rules [3, 31], but where these formulations emphasize probability-of-success objectives, our objective is additive and the induced action rate is a unit-free object determined solely by distributional shape.

On the applied side, existing work on discretionary behavior, whether in the street-level bureaucracy tradition [6, 25] or in the literature on expert responses to algorithmic advice [1, 11, 14, 20], has largely treated each override as a static cost-benefit decision, and does not address the intertemporal rationing problem that arises when overrides are a finite resource. The same gap exists in the growing literature on human-AI decision-making, which studies how experts respond to algorithmic recommendations but not how they manage a finite capacity to deviate from them. From the caseworker's perspective, overriding a VI-SPDAT score [29] and overriding an algorithmic recommendation pose the same intertemporal tradeoff, and the results we develop here apply to both. Our model formalizes the scarcity mechanism emphasized by classic accounts of street-level bureaucracy [4, 25], aligning with service-operations work where limited capacity induces threshold-like upgrading rules [38] and with recent arguments that institutional context is central to the fairness of AI-driven allocation beyond automation-centric framings [32, 34].

Contributions: **(1)** We formalize discretionary overrides of algorithmic (or other scoring-based) recommendations as a finite-horizon dynamic program in which override capacity is a scarce resource, yielding an optimal threshold policy indexed by time remaining and budget remaining. **(2)** We prove that for location-scale improvement distributions, the override rate at every state depends only on distributional shape – not scale or location. Fat-tailed gains induce patient, outlier-targeting behavior; thin-tailed gains induce routine spending. **(3)** Using HMIS data, we recover a baseline algorithmic policy via decision trees and show that observed override patterns shift with capacity openings, consistent with dynamic rationing of a finite discretionary budget, and are front-loaded in the calendar.

2 Model and problem formulation

In each period $t \in \{1, \ldots, T\}$, a single indivisible resource must be assigned to one of two arriving households. Each household $i \in \{1, 2\}$ has a baseline welfare $w_i(t) \overset{\text{iid}}{\sim} \mathcal{D}_w$ and a potential welfare improvement $I_i(t) \overset{\text{iid}}{\sim} I_\theta$ if it receives the resource; all draws are mutually independent. The final welfare of a household receiving the unit is $w_i'(t) = w_i(t) + I_i(t)$. For any period, we use m and M to denote the indices of the households with the **m**inimum and **M**aximum baseline welfare respectively: $m = \arg\min_{i\in\{1,2\}}\{w_i(t)\}$, and $M = \arg\max_{i\in\{1,2\}}\{w_i(t)\}$.

The caseworker actions The *default policy* is a strictly needs-based rule that assigns the resource to the most vulnerable household (the one with the lower baseline welfare), reflecting standard

prioritization logic in services such as homelessness or organ allocation [16, 43]. In contrast, a human decision-maker must balance competing priorities. For example, in homelessness, while the administrative rule is to prioritize the most vulnerable, evaluation is often in terms of success in keeping households stably housed in the future [23], which corresponds to prioritizing the household with a higher increase in underlying welfare level. We thus model the decision-maker as being able to *use discretion* in order to increase total welfare.

2.1 The gain from discretion

The gain from discretion is the welfare difference between the two actions:

$$\Delta_t = (I_M(t) - I_m(t))^+, \tag{1}$$

where m and M index the households with minimum and maximum baseline welfare, respectively (see Appendix C for the full derivation). The gain is non-negative by construction: $\Delta_t > 0$ exactly when the less vulnerable household would benefit more from the resource – an event we call *misalignment*, denoted by A_t^c, with complement A_t (alignment). Under i.i.d. continuous improvements, symmetry gives $p = \Pr(\mathsf{A}_t) = 1/2$ (Remark 3.1 discusses N households, where $p = 1/N$). Across periods, the realized gain follows a mixture:

$$\Delta_t = \begin{cases} 0, & \text{with probability } p \quad \text{(alignment)}, \\ \overset{\text{iid}}{\sim} f_\Delta & \text{with probability } 1-p \quad \text{(misalignment)}, \end{cases} \tag{2}$$

where f_Δ is the density of the gain conditional on misalignment, with conditional CDF $F_{\Delta|\mathsf{A}_t^c}$. We write F_Δ for the unconditional CDF of Δ_t, so that $F_\Delta(x) = p + (1-p)F_{\Delta|\mathsf{A}_t^c}(x)$ for $x \geq 0$. Let $\mathrm{U}_t^{\text{pol}}$ and $\mathrm{U}_t^{\text{max}}$ denote the period-t welfare under the default and welfare-maximizing allocations. By construction $\mathrm{U}_t^{\text{max}} = \mathrm{U}_t^{\text{pol}} + \Delta_t$. Define $\mu_{\text{pol}} = \mathbb{E}[\mathrm{U}_t^{\text{pol}}]$ and $\mu_{\text{max}} = \mathbb{E}[\mathrm{U}_t^{\text{max}}] = \mu_{\text{pol}} + \mathbb{E}[\Delta_t]$ with $\mathbb{E}[\Delta_t] = (1-p)\mathbb{E}[\Delta_t \mid \mathsf{A}_t^c]$ (we delegate these derivations to Appendix C).

2.2 The dynamic program

The goal of the agent is to find an optimal override policy π^* that maximizes total expected welfare over T periods subject to a budget of K overrides:

$$\max_{(d_1,\ldots,d_T)} \sum_{t=1}^{T} \mathbb{E}[U_t(d_t)] \quad \text{s.t.} \quad \sum_{t=1}^{T} d_t \leq K,$$

where $d_t \in \{0,1\}$ indicates whether an override is used at time t. Let the state be (τ, k), with $\tau \in \{0,\ldots,T\}$ periods remaining and $k \in \{0,\ldots,K\}$ budget units left, and let $V(\tau,k)$ denote the maximum expected total welfare from state (τ,k).

Boundary conditions: (i) No time remaining: $V(0,k) = 0$ for all k. (ii) No budget remaining: $V(\tau,0) = \tau \cdot \mu_{\text{pol}}$ for all $\tau \geq 1$; the agent collects the default-policy utility in every remaining period. (iii) Ample budget: $V(\tau,k) = \tau \cdot \mu_{\text{max}}$ for all $k \geq \tau$; the agent can override at every misalignment.

The recursion: Conditioning on alignment or misalignment in the current period and simplifying (see Appendix C.5 for the full algebra), the value function satisfies:

$$V(\tau,k) = \mu_{\text{pol}} + pV(\tau-1,k) + (1-p)\mathbb{E}_{\Delta|\mathsf{A}_t^c}\Big[\max\big\{V(\tau-1,k), \Delta + V(\tau-1,k-1)\big\}\Big] \tag{3}$$

The interpretation is direct: the agent always collects the baseline utility μ_{pol}. If a misalignment occurs (with probability $1-p$), the agent captures the additional gain Δ only if it exceeds the opportunity cost of consuming a unit of budget.

The optimal policy: The maximization in Equation (3) implies that the optimal policy is a state-dependent threshold rule:

$$d_t^*(\tau,k) = \begin{cases} 1 & \text{if } \mathsf{A}_t^c \text{ and } \Delta_t \geq \mathcal{T}_{\tau,k}, \\ 0 & \text{otherwise}, \end{cases}$$

where the threshold

$$\mathcal{T}_{\tau,k} = V(\tau-1,k) - V(\tau-1,k-1) \tag{4}$$

is the marginal continuation value of a budget unit for $\tau \geq 1$, $k \geq 1$. The agent overrides if and only if the current gain exceeds the value of carrying that unit of budget into the future. In the minimal case $T = 2$, $K = 1$, the threshold reduces to $\mathcal{T}_{2,1} = \mathbb{E}[\Delta] = (1-p)\mathbb{E}[\Delta \mid \mathsf{A}_t^c]$: the expected gain from having the override available in the final period (more on this in Section 3.3 and Appendix E).

3 Behavioral invariance

What determines the *rate* at which the optimal policy exercises overrides? The threshold $\mathcal{T}_{\tau,k}$ depends on the units of welfare, but observable behavior is the *probability* that the gain clears it.

3.1 Assumption and base quantities

Assume per-household improvements follow a location-scale family [5] of distributions: $I = a + sX$, where $a \in \mathbb{R}$ is a location parameter, $s > 0$ is a scale parameter, and $X \sim \mathcal{D}$ is a base shape with CDF $F_{\mathcal{D}}$, PDF $f_{\mathcal{D}}$ with non-negative support, and finite mean. Intuitively, a location-scale family is one in which every member can be obtained from a common base distribution by shifting (a) and stretching (s). Examples include the Exponential and Half-Normal families, and the Pareto, Gamma, Weibull, and Lognormal families when their shape parameters are held fixed.

Let $X, X' \overset{\text{iid}}{\sim} \mathcal{D}$ be independent draws. Then, we define two quantities that depend only on $\mathcal{D}$ that we will use in our theorem. First, the **base threshold** (c) which is the expected positive difference between two draws $c(\mathcal{D}) = \mathbb{E}[(X'-X)^+]$, and second, the **difference CDF** (G) i.e., the CDF of the base difference $\Delta_0 = X' - X$:

$$G_{\mathcal{D}}(x) = \Pr(\Delta_0 \leq x) = \int f_{\mathcal{D}}(u) F_{\mathcal{D}}(u+x)\,du. \tag{5}$$

We first give a closed form for $c(\mathcal{D})$ that depends only on the CDF of the base shape.

Proposition 3.1. *$c(\mathcal{D})$ admits a convenient closed form. Suppose $X \sim \mathcal{D}$ has non-negative support and finite mean. Then the base threshold is:*

$$c(\mathcal{D}) = \int_0^\infty F_{\mathcal{D}}(x)\big(1 - F_{\mathcal{D}}(x)\big)\,dx. \tag{6}$$

Sketch of proof. A full derivation is in Appendix D.1. Starting from $c(\mathcal{D}) = \mathbb{E}[(X'-X)^+]$ with $X, X' \overset{\text{iid}}{\sim} \mathcal{D}$, we condition on $X = x$ and write $\mathbb{E}[(X'-X)^+ \mid X = x] = \int_x^\infty (y-x) f_{\mathcal{D}}(y)\,dy$. Integration by parts rewrites this in terms of the tail CDF $1 - F_{\mathcal{D}}(y)$, and we swap the order of integration. The resulting double integral collapses to the single CDF-only expression in (6). The restriction to $[0,\infty)$ uses the non-negative support of improvements. □

3.2 The discretionary value and general invariance

The recursion in Equation (3) mixes two sources of value: the baseline utility μ_{pol} that accrues every period regardless of whether discretion is exercised, and the incremental value that arises from having override budget available. The key to our main result is isolating the part of the value function attributable to having override budget. Define the *discretionary value*:

$$W(\tau,k) = V(\tau,k) - \tau \cdot \mu_{\text{pol}},$$

which measures the expected welfare gain from holding k overrides with τ periods remaining, minus the baseline which is collected at each step. By construction, $W(0,k) = 0$ and $W(\tau,0) = 0$. Substituting into the recursion (3), the constant $\tau \cdot \mu_{\text{pol}}$ cancels from both sides, yielding:

$$W(\tau,k) = p \cdot W(\tau-1,k) + (1-p) \cdot \mathbb{E}_{\Delta \mid \mathsf{A}_t^c}\left[\max\left\{W(\tau-1,k), \Delta + W(\tau-1,k-1)\right\}\right]. \tag{7}$$

Under the location–scale model, the gain from discretion factors as

$$\Delta = (I_M - I_m)^+ = s \cdot (X_M - X_m)^+ = s \cdot \Delta_0^+, \tag{8}$$

where $\Delta_0 = X' - X$ (Appendices D.2.1 and D.2.2 give the algebra for Equations (7) and (8)). The location a cancels in the difference; the scale s factors out because $s > 0$ preserves sign. We are now ready to state our main result.

Theorem 3.1 (General behavioral invariance)**.** *Let per-household improvements follow a location-scale family $I = a + sX$ with $s > 0$, where the base shape $\mathcal{D}$ has non-negative support and finite mean. Then for every state (τ, k) with $\tau \geq 1$ and $k \geq 1$:*

(i) ***The discretionary value scales linearly with s and is invariant to a:*** *there exists a quantity $w(\tau, k; \mathcal{D})$, depending only on $\mathcal{D}$, p, and (τ, k), such that $W(\tau, k) = s \cdot w(\tau, k; \mathcal{D})$.*

(ii) ***The optimal threshold scales linearly with s and is invariant to a:*** *$\mathcal{T}_{\tau,k} = s \cdot \theta_{\tau,k}(\mathcal{D})$, where $\theta_{\tau,k}(\mathcal{D}) = w(\tau - 1, k; \mathcal{D}) - w(\tau - 1, k - 1; \mathcal{D})$ is the base threshold.*

(iii) ***The spending probability at every state is scale- and location-free:*** *$q_{\tau,k} = 1 - G_{\mathcal{D}}(\theta_{\tau,k}(\mathcal{D}))$, where $G_{\mathcal{D}}$ is the CDF of the base difference $\Delta_0 = X' - X$ as defined in* (5)*.*

Proof sketch (full proof in Appendix D.2). **Part (i)** proceeds by induction on τ. The base case $W(0, k) = 0 = s \cdot 0$ is immediate. For the inductive step, assume $W(\tau - 1, j) = s \cdot w(\tau - 1, j; \mathcal{D})$ for all j. Substituting this and $\Delta = s \cdot \Delta_0^+$ into the W-recursion (7), the factor $s > 0$ appears in every term inside the $\max$ and factors out (since multiplying by a positive constant preserves ordering). Pulling s outside the expectation by linearity gives $W(\tau, k) = s \cdot [\cdots]$, where the bracket depends only on $\mathcal{D}$, p, and the base values $w(\tau - 1, \cdot; \mathcal{D})$. **Part (ii)** follows from part (i): $\mathcal{T}_{\tau,k} = W(\tau-1, k)-W(\tau-1, k-1) = s\cdot[w(\tau-1, k; \mathcal{D})-w(\tau-1, k-1; \mathcal{D})] = s\cdot\theta_{\tau,k}(\mathcal{D})$. **Part (iii)** uses both preceding parts. The unconditional CDF of $\Delta = s \cdot \Delta_0^+$ satisfies $F_\Delta(x) = G_{\mathcal{D}}(x/s)$ for $x \geq 0$, since $\Pr(s \cdot \Delta_0^+ \leq x) = \Pr(\Delta_0^+ \leq x/s) = \Pr(\Delta_0 \leq x/s) = G_{\mathcal{D}}(x/s)$, where the second equality holds because $\{\Delta_0^+ \leq x/s\} = \{\Delta_0 \leq x/s\}$ for $x/s \geq 0$. Evaluating at the threshold: $F_\Delta(\mathcal{T}_{\tau,k}) = G_{\mathcal{D}}(s \cdot \theta_{\tau,k}/s) = G_{\mathcal{D}}(\theta_{\tau,k})$, so $q_{\tau,k} = 1 - G_{\mathcal{D}}(\theta_{\tau,k})$. □

Remark 3.1 (Beyond two households and a fixed budget). The proof of Theorem 3.1 uses only two facts about the gain: it factors as $\Delta = s \cdot Z$, where the distribution of Z does not depend on a or s, and gains are independent across periods. Both facts are robust to two natural extensions. If N households arrive in each period, the default serves the neediest household and an override serves the household with the largest improvement, so $\Delta = s \cdot (\max_i X_i - X_m)$, and $p = 1/N$ when improvements are i.i.d. and continuous. The same holds if baselines and improvements are dependent within a period, as long as their joint distribution does not depend on payoff units; p is then whatever that distribution implies. If instead the budget is replenished each period by a random number of units, drawn independently of the welfare draws, the recursion (7) averages the continuation values over the replenishment, and s still factors out of every term. In both cases the argument of Theorem 3.1 carries over step by step. We leave the formal treatment of this to future work.

3.3 A closer look at a non-trivial two-period case

Theorem 3.1 establishes that override rates are shape-determined at every state of the general (T, K) problem. Here, we look at the minimal non-trivial case $(T = 2, K = 1)$ where the full behavioral profile collapses to a single scalar. This case is worth developing in detail for two main reasons. First, it isolates the core spend-or-save tension in its purest form: the agent faces exactly one opportunity to override before the final period, making the tradeoff between immediate use and option value maximally transparent. Second, it admits closed-form expressions that yield concrete, interpretable predictions about when agents front-load versus back-load their discretion.

Theorem 3.2 (Shape dependence of optimal policy, two-period)**.** *Let the improvement I follow a location–scale family $I = a + sX$ with $s > 0$, where the base shape $\mathcal{D}$ has non-negative support and finite mean. Then, when $T = 2$ and $K = 1$:*

(i) ***Optimal threshold scales linearly with s.*** *The two-period threshold $\mathcal{T}_{2,1} = \mathbb{E}[\Delta]$ is invariant to location a and satisfies $\mathcal{T}_{2,1} = s \cdot c(\mathcal{D})$, where the base threshold $c(\mathcal{D})$ depends only on the shape and is given by* (6)*.*

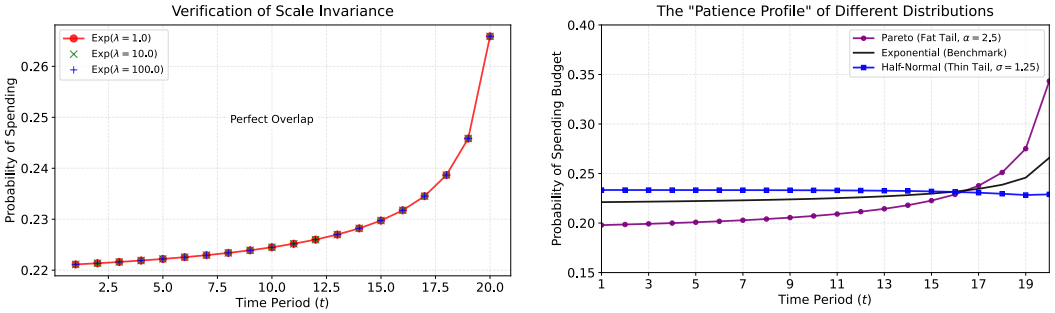
(ii) ***Policy behavior is summarized by a single scalar.*** *All optimal spending probabilities are functions of $\psi(\mathcal{D}) = F_\Delta(\mathcal{T}_{2,1})$, which is invariant to a and s and equals $\psi(\mathcal{D}) = G_{\mathcal{D}}(c(\mathcal{D}))$.*

This theorem can be proved in two ways: directly from the model assumptions (Appendix E.5) or as a special case of Theorem 3.1 (Appendix E.6).

Behavioral implications of shape: The invariance result yields a sharp behavioral dichotomy driven by distributional shape. In our numerical examples, heavier-tailed distributions tend to yield larger ψ. When ψ is sufficiently large, the agent is patient: conservative relative to a full-information oracle in period 1 and back-loading spending to period 2. Smaller values of ψ favor more frequent early overrides. In Appendix E.8, we formalize these comparisons through two lemmas that identify the exact thresholds at which the DP becomes internally patient (Lemma E.2) and more conservative than the oracle (Lemma E.1). We also state a corollary (Corollary E.1) connecting tail thickness to patience for families in which ψ increases with tail thickness.

Numerical verification: Monte Carlo simulations ($N = 10^6$ draws per configuration) confirm both claims of Theorem 3.2; details and closed forms for the Exponential and Half-Normal cases are in Appendix E.7. Scale invariance holds exactly: for every family tested (Exponential, Half-Normal, Gamma, Lognormal), varying the scale parameter 5- to 100-fold leaves $\hat{\psi}$ unchanged to three decimal places (Table 3). Shape, by contrast, moves behavior decisively. Within the Weibull family (mean-normalized to $\mathbb{E}[X] = 1$), moving from a lighter tail ($c = 1.5$, $\hat{\psi} = 0.668$) to a heavier tail ($c = 0.8$, $\hat{\psi} = 0.721$) reduces first-period spending by 16%; Lognormal exhibits an even wider range, from $\hat{\psi} = 0.678$ ($\sigma = 0.5$) to $\hat{\psi} = 0.802$ ($\sigma = 1.5$).

In the extreme configurations of Table 5, estimated patience is high: Lognormal with $\sigma = 5$ yields $\hat{\psi} = 0.994$, and Pareto with $b = 1.05$ yields $\hat{\psi} = 0.962$. At the other extreme, Gamma with $k = 10^4$ (near-degenerate) gives $\hat{\psi} \approx 0.655$, the only example below the patience cutoff of $2/3$ (Corollary E.1): this agent spends slightly more often in period 1 than in period 2. Tables 3– 5 in Appendix E.7.3 report the complete results.



(a) The optimal spending trajectory for an Exponential distribution remains identical despite varying the scale parameter λ (the inverse of the scale) by two orders of magnitude, confirming Theorem 3.1 (iii).

(b) Spending probability over time for distributions with different tail shapes. Heavy-tailed distributions (Pareto) induce a patient policy; thin-tailed distributions (Half-Normal) spend at a nearly constant rate.

Figure 1: Scale invariance and shape dependence ($T = 20, K = 5$): The optimal spending probability is invariant to the scale of rewards (a) and is driven by the tail of the gain distribution (b). The two panels use different vertical scales.

3.4 Scale invariance and shape dependence in the general problem

Since every $q_{\tau,k}$ is scale- and location-free (Theorem 3.1), and the budget distribution $\pi_t(k)$ is computed from these probabilities via a deterministic forward pass (see Appendix F), the entire observable spending profile inherits the same invariance. Figure 1 verifies this numerically and illustrates the behavioral consequences. This is the same dichotomy as demonstrated in the $(T = 2, K = 1)$ case. In our fat-tailed example, most override opportunities yield modest improvements but rare cases produce very large gains. The option value of waiting is high, so the optimal agent is patient – conserving overrides for outlier cases and concentrating spending in the final periods as the value of saving an override declines. On the other hand, in our thin-tailed example, waiting is less likely to reveal a dramatically better opportunity, and the agent overrides earlier and more routinely.

This pattern is visible in Figure 1b: Pareto gains (heavy tail) produce a sharply back-loaded spending curve, while Half-Normal gains (thin tail) produce a nearly flat one.

4 Empirical evidence from homeless services

To complement our theoretical characterization of threshold-based override behavior, we examine whether discretionary overrides in a real allocation system vary with the operational constraints that our model identifies as drivers of threshold-based spending. The setting is homeless-service placement in St. Louis, where caseworkers assign households to either Emergency Shelter (ES, the lower-intensity default) or Transitional Housing (TH, the scarce, resource-intensive intervention). Because the latent gains Δ are not directly observed, our goal is not to estimate model primitives but to test whether the *correlates* of overrides are consistent with the dynamic rationing behavior our model predicts. We ask two questions. First, how does discretionary behavior respond to shifts in the operational constraints that bind the caseworker? And second, how is the use of discretion spread over the calendar?

Setting and discretion measurement: We analyze administrative records from the HMIS covering 2008–2014, when caseworkers relied on professional judgment rather than formulaic prioritization tools. Following Pokharel et al. [33], we define a baseline "policy" by training a short, interpretable decision tree (SER-DT [39], constrained to a depth of 4) on 34 household features. The tree's prediction P_i represents the dominant assignment heuristic; a *discretionary event* occurs when the caseworker's actual placement differs from this prediction ($A_i \neq P_i$). Appendix G.3 discusses how errors in this recovered baseline could affect our estimates, and Appendix G.7 shows that the estimates in Table 1 stay within their 95% confidence intervals when tree depth ranges from 2 to 6.

We distinguish *upgrading* (P_i=ES, A_i=TH: spending the budget) from *rationing* (P_i=TH, A_i=ES: conserving it). Whereas the earlier study asks which households receive overrides, we ask how override use varies over time. This temporal analysis is new to this paper: we aggregate overrides by day and construct measures of the system's operating conditions, including recent exits, to examine how discretion varies with capacity and the calendar. The theoretical model remains our main contribution, and these results offer an initial check of its predictions in one setting.

Specification: We aggregate to the daily level and model override counts as binomial draws, with a logit link and covariates capturing the system's operational state. The key operational regressors are rolling 7-day ES and TH assignments (proxying workload and throughput burden) and rolling 7-day ES and TH exits (proxying short-run capacity *openings*). In the language of the model, exits serve as a reduced-form proxy for the tightness of intervention-specific capacity: when exits are high, slots have recently been freed and the effective shadow price of that intervention falls; when exits are low, the intervention is congested and overrides into it are costlier. In the model, this shadow price is the threshold $\mathcal{T}_{\tau,k}$ in Equation (4). Recent exits do not change p, the probability that the default already matches the caseworker's judgment, which depends on the mix of cases; they lower the threshold. Since the threshold falls as the remaining budget grows (Figure 2 in Appendix F.3), an opening has the same effect as a larger budget. We also include a lagged TH placement share, day-type indicators (Monday / Tue–Fri / Weekend), month-of-year fixed effects (October reference), and a federal-holiday indicator. All operational features use strictly lagged information to avoid mechanical dependence on same-day decisions. For directional models, denominators are restricted to the eligible trial set (e.g., upgrading is only possible when P_i=ES). The full specification and a weekly-block robustness check appear in Appendix G.

4.1 Results

Table 1 reports selected odds ratios from the rolling specification; Appendix H.1 has the full tables.

Capacity openings and substitution: The rationing direction shows the most consistent association with capacity constraints. n Table 1, each additional ES exit in the prior week (freeing shelter beds) is associated with more rationing (OR $= 1.026$, $p < 0.01$), while each additional TH exit (freeing transitional slots) is associated with less rationing (OR $= 0.952$, $p < 0.001$). This sign pattern matches a substitution logic. When TH slack increases, the need to conserve TH falls; when ES slack increases, defaulting households down to ES becomes easier. For upgrading, the pattern is complementary. Upgrades decline when ES exits are high (OR $= 0.971$, $p < 0.01$) and rise

Table 1: Summary of correlates of daily discretion (rolling specification; day-type baseline Tue–Fri).

Predictor	Overall (p_t^{all}) OR	Overall [95% CI]	Upgrading $(p_t^{\uparrow})$ OR	Upgrading [95% CI]	Rationing $(p_t^{\downarrow})$ OR	Rationing [95% CI]
Operational state (prior 7 days $[t-7, t-1]$)						
THShare_{t-1}	1.060	[0.859, 1.308]	0.931	[0.683, 1.271]	1.150	[0.853, 1.551]
ES exits ($[t-7, t-1]$)	0.995	[0.982, 1.008]	0.971**	[0.953, 0.989]	1.026**	[1.007, 1.045]
TH exits ($[t-7, t-1]$)	0.995	[0.978, 1.012]	$1.021^{\dagger}$	[0.998, 1.045]	0.952***	[0.928, 0.977]
Timing (reference: Tue–Fri)						
Mon	1.244**	[1.074, 1.440]	1.140	[0.919, 1.415]	1.241*	[1.006, 1.531]
Weekend	0.762*	[0.605, 0.961]	0.813	[0.574, 1.153]	0.639**	[0.465, 0.879]
Seasonality (reference: October)						
Nov	0.705*	[0.512, 0.970]	0.880	[0.544, 1.424]	$0.677^{\dagger}$	[0.427, 1.075]
Dec	0.592**	[0.428, 0.818]	0.730	[0.448, 1.190]	0.547*	[0.345, 0.869]
Jan	0.789	[0.585, 1.063]	1.125	[0.720, 1.758]	0.612*	[0.398, 0.943]
Jun	0.710*	[0.524, 0.962]	1.027	[0.647, 1.630]	0.522**	[0.339, 0.804]

Notes. Entries are odds ratios from binomial-logit GLMs estimated on daily aggregates. Rolling operational features use the prior-week window $[t-7, t-1]$. All models include the full set of month-of-year indicators (October reference) as well as a federal-holiday indicator; the table reports selected coefficients for readability. Full coefficient tables, the weekly block specification, and an alternative day-of-week fixed-effects specification are reported in Appendices H.1–H.2. Stars indicate significance: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$, $^{\dagger}p < 0.1$.

(marginally) when TH exits are high (OR $= 1.021$, $p < 0.1$). These odds ratios are per exit, and the number of exits in a week can vary, so the effect of a typical change is larger. Per standard deviation of the prior-week exit count (5.7 ES exits and 4.4 TH exits), ES exits are associated with 16% higher odds of rationing and TH exits with 19% lower odds (OR $= 1.157$ and 0.807), and with 16% lower and 10% higher odds of upgrading (OR $= 0.843$ and 1.097); the p-values are those of the per-exit estimates. These are reported in Tables 8 and 9 in Appendix H. These associations are conditional on the baseline recommendation, so they do not merely reflect mechanical availability effects. When slots are plentiful, the effective cost of diverting one to an override is lower, which in the model corresponds to a lower threshold for spending an override. Pokharel et al. [33] found that caseworkers in this system are more likely to exercise discretion on lower-acuity households, which suggests a complementary channel: both higher slot availability and lower case acuity reduce the professional scrutiny attached to deviating from the default. On its own, the capacity association does not separate our model from the idea that discretion responds to scarcity, so we treat it as a consistency check.

Temporal pattern of discretion: Discretion is front-loaded across multiple calendar cycles. Within the week, overall discretion is elevated on Mondays relative to typical weekdays (OR $= 1.244$, $p < 0.01$). At the fiscal-year level, relative to October (the federal fiscal-year start), overall discretion is modestly lower in most subsequent months, with the sharpest reductions in November-December and June. This temporal pattern i.e., higher override rates at the beginning of the relevant calendar cycle is what our theory associates with thin-tailed improvement distributions, where the option value of waiting is modest and routine early spending is optimal.

Discretion is also lower on weekends, overall (OR $= 0.762$, $p < 0.05$) and especially for rationing (OR $= 0.639$, $p < 0.01$; Appendix H.1). However, we note that operational schedules could also produce this weekly pattern. Fewer staff work on weekends and intake slows, so weekend placements may stick closer to the default, while decisions that need more attention wait until Monday, when the weekend's cases are processed. The fiscal year could work the same way: October starts the federal fiscal year, when budgets and reporting cycles reset. Staffing and intake limits are themselves a form of the operational scarcity the model describes. These operational constraints could make Mondays busier, but volume by itself cannot raise our outcome, which is the share of a day's placements that override the baseline; more cases raise the denominator along with the numerator. What staffing and intake schedules can change, however, is the mix of cases. A household that waits out the weekend may arrive on Monday in worse shape than a typical weekday case, so Monday's pool may contain more cases in which judgment disagrees with the baseline, and that alone would raise the override rate without any budget rationing. Our data cannot separate these explanations from the

model's prediction that thin-tailed gains produce front-loaded spending (see also Appendix H.3), so the calendar results should be read as a description of when discretion is used.

Taken together, two patterns stand out. First, capacity openings redirect discretion: when recent exits free slots in one intervention, caseworkers override toward that intervention more readily, consistent with a lower threshold for spending. Second, discretion is front-loaded in the calendar, peaking on Mondays within each week and elevated early in the fiscal year. This temporal pattern resembles the routine early spending our theory associates with thin-tailed improvement distributions, where waiting is unlikely to reveal a dramatically better case. The overall picture is one in which discretion behaves like a scarce, dynamically managed resource, and caseworkers exercise it more freely when conditions reduce the cost of doing so.

5 Discussion and conclusion

Frontline public-service workers have always rationed their judgment, but formal models of how that rationing should unfold over time have been largely absent. We address this by modeling discretion as a dynamic allocation problem in which overrides of an algorithmic default are a finite resource: spending one today reduces what is available tomorrow. The optimal policy is a state-dependent threshold rule, and the main theoretical result is a behavioral invariance. For location-scale families of improvement distributions, the rate at which an optimal agent exercises discretion is independent of payoff units and determined entirely by the shape of the gain distribution. Because this invariance holds at every state of the (T, K) dynamic program, the entire observable spending profile is a unit-free object, determined by distributional shape, the alignment probability, and the problem dimensions. Fat-tailed gains produce patient, outlier-targeting agents; thin-tailed gains produce routine, front-loaded spenders. In homelessness services, overrides shift toward interventions whose slots have recently opened, as a lower threshold for spending would imply; we also find some evidence suggesting that discretion is front-loaded in the week and the fiscal year.

These results rest on modeling choices that isolate the spend-or-save tradeoff in its sharpest form: i.i.d. gains from a location-scale family of improvements, and a hard budget constraint. Real-world discretion can additionally be shaped by evolving caseload, soft institutional constraints, and learning. Relaxing these assumptions is a natural next step. With a noisy signal of the gain, the optimal rule could keep its threshold form, applied to the expected gain given the signal. There, we expect the invariance to hold when the noise scales with the gain, but it is unclear when the noise has a fixed size. Furthermore, our data show only the front-loaded regime; settings where a rare override can prevent a disaster, such as intensive care triage, should show the patient one, and are worth investigating. Finally, the shape of the gain distribution could also be estimated from downstream outcomes, correcting for the fact that overrides are observed only above the threshold. This paper provides the framework that such extensions would build on.

These results have practical implications for how allocation systems are designed. Whether the standing default is a scoring instrument or an algorithmic recommendation, current tools focus on flagging *which* cases merit attention. Our model suggests they should also surface *when* to act, making the option value of waiting legible to the caseworker managing a finite override budget. The shape-based characterization clarifies why similar organizations may exhibit different discretionary "patience" even under comparable resources: the answer lies not in the scale of potential gains, which washes out, but in their distributional shape. Any system designed to support effective use of discretion should take these factors into account.

Acknowledgments and Disclosure of Funding

We are grateful for support from NSF Award 2533162. We also thank our various community partners who helped conceptualize the challenges facing the delivery of homeless services, as well as their ongoing efforts to support local families.

References

[1] V. Angelova, W. Dobbie, and C. S. Yang. Algorithmic recommendations and human discretion. *Review of Economic Studies*, Sept. 2025. ISSN 0034-6527, 1467-937X. doi: 10.1093/restud/

rdaf084. URL https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdaf084/8254085.

[2] A. Arlotto and X. Xie. Logarithmic regret in the dynamic and stochastic knapsack problem with equal rewards. *Stochastic Systems*, 10(2):170–191, June 2020. ISSN 1946-5238. doi: 10.1287/stsy.2019.0055. URL https://pubsonline.informs.org/doi/10.1287/stsy.2019.0055.

[3] F. T. Bruss. Sum the odds to one and stop. *The Annals of Probability*, 28 (3):1384–1391, June 2000. ISSN 0091-1798. doi: 10.1214/aop/1019160340. URL https://projecteuclid.org/journals/annals-of-probability/volume-28/issue-3/Sum-the-odds-to-one-and-stop/10.1214/aop/1019160340.full.

[4] P. C. Butler, A. Bowers, A. P. Smith, S. R. Cohen-Hatton, and R. C. Honey. Decision making within and outside standard operating procedures: Paradoxical use of operational discretion in firefighters. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 65(7): 1422–1434, Nov. 2023. ISSN 0018-7208, 1547-8181. doi: 10.1177/00187208211041860. URL https://journals.sagepub.com/doi/10.1177/00187208211041860.

[5] G. Casella and R. L. Berger. *Statistical Inference*, chapter 3, pages 116–121. Duxbury, Pacific Grove, CA, 2nd edition, 2002. Section 3.5: Location and Scale Families.

[6] A. Chang. A formal model of street-level bureaucracy. *Rationality and Society*, 34(1):6–27, Feb. 2022. ISSN 1043-4631, 1461-7358. doi: 10.1177/10434631211043205. URL https://journals.sagepub.com/doi/10.1177/10434631211043205.

[7] T. Chen and C. Guestrin. XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 785–794, San Francisco, California, USA, Aug. 2016. ACM. ISBN 9781450342322. doi: 10.1145/2939672.2939785. URL https://dl.acm.org/doi/10.1145/2939672.2939785.

[8] H.-F. Cheng, L. Stapleton, A. Kawakami, V. Sivaraman, Y. Cheng, D. Qing, A. Perer, K. Holstein, Z. S. Wu, and H. Zhu. How child welfare workers reduce racial disparities in algorithmic decisions. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, CHI '22, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450391573. doi: 10.1145/3491102.3501831. URL https://doi.org/10.1145/3491102.3501831.

[9] M. De-Arteaga, R. Fogliato, and A. Chouldechova. A case for humans-in-the-loop: Decisions in the presence of erroneous algorithmic scores. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, CHI '20, pages 1–12, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450367080. doi: 10.1145/3313831.3376638. URL https://doi.org/10.1145/3313831.3376638.

[10] B. C. Dean, M. X. Goemans, and J. Vondrák. Approximating the stochastic knapsack problem: The benefit of adaptivity. *Mathematics of Operations Research*, 33(4):945–964, Nov. 2008. ISSN 0364-765X, 1526-5471. doi: 10.1287/moor.1080.0330. URL https://pubsonline.informs.org/doi/10.1287/moor.1080.0330.

[11] B. J. Dietvorst, J. P. Simmons, and C. Massey. Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General*, 144 (1):114–126, 2015. ISSN 1939-2222, 0096-3445. doi: 10.1037/xge0000033. URL https://doi.apa.org/doi/10.1037/xge0000033.

[12] J. P. Gilbert and F. Mosteller. Recognizing the maximum of a sequence. *Journal of the American Statistical Association*, 61(313):35–73, 1966. doi: 10.1080/01621459.1966.10502008. URL https://www.tandfonline.com/doi/abs/10.1080/01621459.1966.10502008.

[13] B. Green. The flaws of policies requiring human oversight of government algorithms. *Computer Law & Security Review*, 45:105681, July 2022. ISSN 2212473X. doi: 10.1016/j.clsr.2022.105681. URL https://linkinghub.elsevier.com/retrieve/pii/S0267364922000292.

[14] Z. Guo, Y. Wu, J. D. Hartline, and J. Hullman. A decision theoretic framework for measuring AI reliance. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency*, pages 221–236, Rio de Janeiro, Brazil, June 2024. ACM. ISBN 9798400704505. doi: 10.1145/3630106.3658901. URL https://dl.acm.org/doi/10.1145/3630106.3658901.

[15] Z. Guo, Y. Wu, J. Hartline, and J. Hullman. Explaining and improving information complementarities in multi-agent decision-making. *arXiv preprint arXiv:2502.06152*, 2025. doi: 10.48550/ARXIV.2502.06152. URL https://arxiv.org/abs/2502.06152.

[16] Health Resources and Services Administration (HRSA). Questions and answers about liver allocation, Dec. 2025. URL https://www.hrsa.gov/optn/patients/resources/liver/liver-allocation-questions.

[17] T. Hill and R. Kertz. A survey of prophet inequalities in optimal stopping theory. *Contemporary Mathematics*, 125:191–207, Jan. 1992. URL https://digitalcommons.calpoly.edu/rgp_rsr/6.

[18] J. Jiang, W. Ma, and J. Zhang. Tight guarantees for multiunit prophet inequalities and online stochastic knapsack. *Operations Research*, 73(3):1703–1721, May 2025. ISSN 0030-364X, 1526-5463. doi: 10.1287/opre.2022.0309. URL https://pubsonline.informs.org/doi/10.1287/opre.2022.0309.

[19] D. P. Kennedy and R. P. Kertz. Limit theorems for threshold-stopped random variables with applications to optimal stopping. *Advances in Applied Probability*, 22(2):396–411, June 1990. doi: 10.2307/1427542.

[20] S. Kesavan and T. Kushwaha. Field experiment on the profit implications of merchants' discretionary power to override data-driven decision-making tools. *Management Science*, 66 (11):5182–5190, Nov. 2020. ISSN 0025-1909, 1526-5501. doi: 10.1287/mnsc.2020.3743. URL https://pubsonline.informs.org/doi/10.1287/mnsc.2020.3743.

[21] R. Kleinberg and S. M. Weinberg. Matroid prophet inequalities and applications to multi-dimensional mechanism design. *Games and Economic Behavior*, 113:97–115, Jan. 2019. ISSN 08998256. doi: 10.1016/j.geb.2014.11.002. URL https://linkinghub.elsevier.com/retrieve/pii/S0899825614001651.

[22] A. J. Kleywegt and J. D. Papastavrou. The dynamic and stochastic knapsack problem. *Operations Research*, 46(1):17–35, Feb. 1998. doi: 10.1287/opre.46.1.17. URL https://ideas.repec.org/a/inm/oropre/v46y1998i1p17-35.html.

[23] A. R. Kube, S. Das, and P. J. Fowler. Fair and efficient allocation of scarce resources based on predicted outcomes: implications for homeless service delivery. *Journal of Artificial Intelligence Research*, 76:1219–1245, 2023.

[24] J. Lin, C. Zhong, D. Hu, C. Rudin, and M. Seltzer. Generalized and scalable optimal sparse decision trees. In *Proceedings of the 37th International Conference on Machine Learning*, ICML'20, pages 6150–6160, 2020.

[25] M. Lipsky. *Street-level bureaucracy: Dilemmas of the individual in public services*. Russell Sage Foundation, New York, 1980.

[26] J. J. McCall. Economics of information and job search. *The Quarterly Journal of Economics*, 84(1):113–126, Feb. 1970. ISSN 00335533. doi: 10.2307/1879403. URL https://academic.oup.com/qje/article-lookup/doi/10.2307/1879403.

[27] H. McTavish, C. Zhong, R. Achermann, I. Karimalis, J. Chen, C. Rudin, and M. Seltzer. Fast sparse decision tree optimization via reference ensembles. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(9):9604–9613, June 2022. ISSN 2374-3468, 2159-5399. doi: 10.1609/aaai.v36i9.21194. URL https://ojs.aaai.org/index.php/AAAI/article/view/21194.

[28] NIST. Measures of skewness and kurtosis. NIST/SEMATECH e-Handbook of Statistical Methods, Feb. 2026. URL https://www.itl.nist.gov/div898/handbook/eda/section3/eda35b.htm. Accessed 2026-02-05.

[29] OrgCode Consulting. Vulnerability index-service prioritization decision assistance tool (VI-SPDAT) prescreen triage tool for single adults, 2015.

[30] G. Persad, A. Wertheimer, and E. J. Emanuel. Principles for allocation of scarce medical interventions. *The Lancet*, 373(9661):423–431, Jan. 2009. ISSN 01406736. doi: 10.1016/S0140-6736(09)60137-9. URL https://linkinghub.elsevier.com/retrieve/pii/S0140673609601379.

[31] J. D. Petruccelli. On a best choice problem with partial information. *The Annals of Statistics*, 8 (5):1171–1174, 1980. ISSN 00905364, 21688966. URL http://www.jstor.org/stable/2240449.

[32] G. Pokharel. Beyond automation: Understanding fairness, ethics, and human discretion in AI-driven societal decisions. *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 8(3):2918–2920, Oct. 2025. ISSN 3065-8365. doi: 10.1609/aies.v8i3.36793. URL https://ojs.aaai.org/index.php/AIES/article/view/36793.

[33] G. Pokharel, S. Das, and P. J. Fowler. Discretionary trees: Understanding street-level bureaucracy via machine learning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38 (20):22303–22312, Mar. 2024. ISSN 2374-3468, 2159-5399. doi: 10.1609/aaai.v38i20.30236. URL https://ojs.aaai.org/index.php/AAAI/article/view/30236.

[34] G. Pokharel, S. Farabi, P. J. Fowler, and S. Das. Street-level AI: Are large language models ready for real-world judgments? *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 8(3):2043–2054, Oct. 2025. ISSN 3065-8365. doi: 10.1609/aies.v8i3.36694. URL https://ojs.aaai.org/index.php/AIES/article/view/36694.

[35] J. Preater. On multiple choice secretary problems. *Mathematics of Operations Research*, 19 (3):597–602, Aug. 1994. ISSN 0364-765X, 1526-5471. doi: 10.1287/moor.19.3.597. URL https://pubsonline.informs.org/doi/10.1287/moor.19.3.597.

[36] M. Rothschild and J. E. Stiglitz. Increasing risk: I. A definition. *Journal of Economic Theory*, 2(3):225–243, Sept. 1970. ISSN 00220531. doi: 10.1016/0022-0531(70)90038-4. URL https://linkinghub.elsevier.com/retrieve/pii/0022053170900384.

[37] D. Saxena and S. Guha. Algorithmic harms in child welfare: Uncertainties in practice, organization, and street-level decision-making. *ACM Journal on Responsible Computing*, 1(1), Mar. 2024. doi: 10.1145/3616473. URL https://doi.org/10.1145/3616473.

[38] R. A. Shumsky and F. Zhang. Dynamic capacity management with substitution. *Operations Research*, 57(3):671–684, June 2009. ISSN 0030-364X, 1526-5463. doi: 10.1287/opre.1080.0610. URL https://pubsonline.informs.org/doi/10.1287/opre.1080.0610.

[39] V. F. Souza, F. Cicalese, E. Laber, and M. Molinaro. Decision trees with short explainable rules. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 12365–12379. Curran Associates, Inc., 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/500637d931d4feb99d5cce84af1f53ba-Paper-Conference.pdf.

[40] C. Starke, J. Baleis, B. Keller, and F. Marcinkowski. Fairness perceptions of algorithmic decision-making: A systematic review of the empirical literature. *Big Data & Society*, 9 (2), July 2022. ISSN 2053-9517, 2053-9517. doi: 10.1177/20539517221115189. URL https://journals.sagepub.com/doi/10.1177/20539517221115189.

[41] R. T. Sutton, D. Pincock, D. C. Baumgart, D. C. Sadowski, R. N. Fedorak, and K. I. Kroeker. An overview of clinical decision support systems: benefits, risks, and strategies for success. *npj Digital Medicine*, 3:17, Feb. 2020. ISSN 2398-6352. doi: 10.1038/s41746-020-0221-y. URL https://www.nature.com/articles/s41746-020-0221-y.

[42] M. Tamaki. A secretary problem with double choices. *Journal of the Operations Research Society of Japan*, 22(4):257–265, 1979. ISSN 0453-4514, 2188-8299. doi: 10.15807/jorsj.22.257. URL https://www.jstage.jst.go.jp/article/jorsj/22/4/22_KJ00001202221/_article.

[43] U.S. Department of Housing and Urban Development, Office of Community Planning and Development. Notice establishing additional requirements for a Continuum of Care centralized or coordinated assessment system. Notice CPD-17-01, U.S. Department of Housing and Urban Development, Jan. 2017. URL https://www.hud.gov/sites/documents/17-01cpdn.pdf.

[44] M. L. Weitzman. Optimal search for the best alternative. *Econometrica*, 47(3):641–654, May 1979. ISSN 00129682. doi: 10.2307/1910412. URL https://EconPapers.repec.org/RePEc:ecm:emetrp:v:47:y:1979:i:3:p:641-54.

A Table of notation

Table 2: Summary of notation used throughout the paper.

Symbol	Description	Introduced
Problem Dimensions		
T	Total number of periods (finite horizon)	§1
K	Total override budget (number of discretionary actions)	§1
t	Period index, $t \in \{1, \dots, T\}$	§2
τ	Periods remaining (incl. current), $\tau = T - t + 1$	§2.2
k	Remaining budget, $k \in \{0, 1, \dots, K\}$	§2.2
d_t	Override decision indicator at time t (1 = override, 0 = comply)	§2.2
Households and Welfare		
i	Household index, $i \in \{1, 2\}$	§2
$w_i(t)$	Baseline welfare of household i in period t; $w_i(t) \overset{\text{iid}}{\sim} \mathcal{D}_w$	§2
$I_i(t)$	Welfare improvement for household i if given the resource; $I_i(t) \overset{\text{iid}}{\sim} I_\theta$	§2
m, M	Indices of min and max baseline-welfare households	§2
U^{pol}	Total welfare under the default (needs-based) policy	§2.1
$\mathrm{U}^{\max}$	Total welfare under the welfare-maximizing allocation	§2.1
μ_{pol}	Expected per-period welfare under default policy, $\mathbb{E}[\mathrm{U}^{\text{pol}}]$	§2.1
$\mu_{\max}$	Expected per-period welfare under best allocation, $\mathbb{E}[\mathrm{U}^{\max}]$	§2.1
Alignment and the Gain from Discretion		
A_t	Alignment event: default policy is welfare-maximizing ($I_m \geq I_M$)	§2.1
A_t^c	Misalignment event: override can improve welfare ($I_M > I_m$)	§2.1
p	Alignment probability, $p = \Pr(\mathsf{A}_t)$; equals $1/2$ under i.i.d. continuous improvements	§2.1
Δ_t	Gain from discretion at time t: $\Delta_t = (I_M - I_m)^+$	§2.1
f_Δ	Density of Δ_t conditional on A_t^c	§2.1
F_Δ	Unconditional CDF of Δ_t, including its mass at 0	§2.1
Location–Scale Family and Base Shape		
$I = a + sX$	Location–scale model for improvements	§3.1
a	Location parameter	§3.1
s	Scale parameter ($s > 0$)	§3.1
$\mathcal{D}$	Base shape distribution; $X \sim \mathcal{D}$	§3.1
$F_{\mathcal{D}}, f_{\mathcal{D}}$	CDF and PDF of the base shape $\mathcal{D}$	§3.1
Δ_0	Base difference, $\Delta_0 = X' - X$ for $X, X' \overset{\text{iid}}{\sim} \mathcal{D}$	§3.1
$G_{\mathcal{D}}(x)$	CDF of Δ_0: $G_{\mathcal{D}}(x) = \int f_{\mathcal{D}}(u)\, F_{\mathcal{D}}(u + x)\, du$	Eq. (5)
$c(\mathcal{D})$	Base threshold, $c(\mathcal{D}) = \mathbb{E}[(X' - X)^+]$	§3.1
Value Functions and Thresholds		
$V(\tau, k)$	Optimal expected total welfare from state (τ, k)	§2.2
$W(\tau, k)$	Discretionary value: $W(\tau, k) = V(\tau, k) - \tau \cdot \mu_{\text{pol}}$	§3.2
$w(\tau, k; \mathcal{D})$	Base discretionary value (scale-free): $W(\tau, k) = s \cdot w(\tau, k; \mathcal{D})$	Thm. 3.1
$\mathcal{T}_{\tau,k}$	Optimal override threshold at state (τ, k)	§2.2
$\theta_{\tau,k}(\mathcal{D})$	Base threshold (scale-free): $\mathcal{T}_{\tau,k} = s \cdot \theta_{\tau,k}(\mathcal{D})$	Thm. 3.1
Behavioral Quantities		
ψ	Patience scalar (two-period): $\psi = F_\Delta(\mathcal{T}_{2,1}) = G_{\mathcal{D}}(c(\mathcal{D}))$	Thm. 3.2
$q_{\tau,k}$	Conditional spending prob. at state (τ, k): $q_{\tau,k} = 1 - F_\Delta(\mathcal{T}_{\tau,k})$	§3.4
$\pi_t(k)$	Prob. of having budget k at the start of period t	§3.4
π^*	Optimal override policy	§2.2
$\Pr^{\pi^*}$	Probability under the optimal policy π^*	App. E
$\Pr^{\tilde{\pi}}$	Probability under the oracle (full-information) policy $\tilde{\pi}$	App. E

B Relationship to prior work

Our model of budgeted discretion is connected to a broad literature on sequential decision-making under scarcity. Across these settings, a decision-maker observes noisy, per-period opportunities and

must decide when to convert them into a small number of irreversible actions. A recurring structural feature is a state-dependent threshold: a scarce action is taken only when the realized opportunity exceeds a continuation-value benchmark. While threshold policies are a familiar structural feature of these problems, our contribution is a behavioral invariance result that characterizes how often an optimal policy acts as the scale of payoffs changes. The related work below traces the lineage of this viewpoint, moving from budgeted online allocation benchmarks to the threshold structure of stopping and search, then to comparative statics driven by tail behavior, and finally to invariance arguments that produce scale-free behavioral predictions.

Budgeted online choice and clairvoyant benchmarks. Discretion can be viewed as an online allocation problem with a limited inventory: each period reveals a stochastic *opportunity* (the incremental gain from overriding a default), and the agent decides whether to spend one unit of a finite budget or save it. This abstraction closely parallels dynamic and stochastic knapsack models, where optimal policies compare realized rewards to the shadow value of remaining capacity [22]. Related work in stochastic packing emphasizes the effectiveness of simple adaptive policies under uncertainty [10, 2]. A complementary benchmark tradition compares online decisions to an offline "prophet" who observes all realizations in advance [17, 21, 18]. In contrast to this value-focused literature, we use the finite-horizon optimum to characterize *the rate* at which discretionary capacity is consumed. This behavioral lens naturally leads to the optimal-stopping and search tradition, where thresholds are interpreted as reservation values balancing current gains against the option value of waiting.

Threshold structure with limited selections. The tension between acting now and waiting for future opportunities is central to optimal stopping and search. In canonical job search models, a decision-maker accepts an offer only if it exceeds a reservation value that equates the marginal benefit of stopping to the option value of continued search [26, 44]. Closely related formulations appear in best-choice and secretary problems, where a sequence of realizations is observed and decisions are irrevocable; in full-information variants, optimal behavior is governed by time-varying thresholds that become less selective as the horizon shortens [12]. Multi-choice extensions allow multiple acceptances and yield thresholds indexed by remaining selections, mirroring the (τ, k) state of our model [42, 35]. While these problems are often framed in terms of selecting the best realization or maximizing expected selected value, we use the same threshold structure to study how frequently an optimal agent exercises a scarce action when decisions are coupled through a finite budget. This framing also highlights that the shape of the opportunity distribution matters: the value of waiting depends on how much upside remains in the tail.

Tail geometry and selective action. A central comparative-statics insight from search and stopping is that greater dispersion can increase selectivity by amplifying the upside of waiting while leaving downside avoidable [36]. In heavy-tailed environments, threshold-based stopping rules can effectively target rare extreme realizations, a phenomenon highlighted in asymptotic analyses of threshold-stopped processes [19]. Our work brings this intuition into a finite-budget setting: patience depends jointly on time and remaining discretionary capacity, and tail heaviness induces systematically more outlier-targeting use of discretion.

Invariance and scale-free behavior. Robust sequential decision rules often arise from exploiting invariances of the underlying observation model. In best-choice problems with partial information, rank-based formulations yield stopping behavior that is insensitive to distributional parameters [31]. Related approaches, such as the odds algorithm, provide compact index rules with transparent comparative statics [3]. Whereas these classic formulations often emphasize probability-of-success objectives (e.g., selecting the overall best realization or the last success), our objective is additive: after accounting for the baseline policy, the agent maximizes the expected sum of incremental gains Δ_t obtained from at most K overrides. Building on this perspective, our main theoretical contribution is a behavioral invariance: for location–scale families of improvement distributions, the *probability* that an optimal agent spends an override at a given state is invariant to affine rescaling of payoffs. While numerical thresholds necessarily depend on units, the induced *action rate* is a unit-free object determined solely by distributional shape, yielding predictions about discretionary behavior across environments. These scale-free predictions are particularly useful in empirical settings, where payoffs are measured imperfectly and discretion is shaped by institutional constraints.

Discretion under organizational scarcity. Empirically, discretionary overrides are pervasive in high-stakes services. Behavioral studies document under-use of algorithmic recommendations following observed errors [11] and strategic overrides to repair perceived biases [37, 8, 14, 15]. Economic field evidence similarly treats discretion as a channel for private information and local judgment [1, 20], but often models overriding as a static choice. Complementary work in public administration has proposed formal models of street-level discretion in which frontline agents choose a level of engagement by trading off public-service benefits against personal costs such as stress, time, and information acquisition, under constraints on resources and authority [6]. These models capture why discretion is costly and institutionally bounded, but do not address the intertemporal rationing problem that arises when overrides are treated as a finite resource. Classic accounts of street-level bureaucracy instead emphasize that discretion is fundamentally shaped by caseload pressure and downstream scarcity [25, 4]. Our model formalizes this scarcity mechanism and aligns with service-operations work in which limited capacity induces protection levels and threshold-like upgrading rules [38]. We use homeless-service allocation as an empirical setting to connect these ideas, showing that observed override behavior tracks capacity in a manner consistent with optimal management of discretion under scarcity [23], and aligning with recent arguments that discretion and institutional context are central to the fairness and ethics of AI-driven societal decisions beyond purely automation-centric framings [32, 34].

Taken together, this literature scaffolds our analysis. We bring these threads together in a setting motivated by discretionary overrides, where the primitive is not choosing the best item but choosing when to deviate from a standing default under a hard budget. Accordingly, our threshold is not calibrated to maximize a success probability, but emerges as the shadow price of discretionary capacity in a welfare-maximizing DP. By focusing on the induced spending probabilities – a behaviorally observable object – we obtain a scale-free characterization of discretion that links the geometry of the gain distribution to systematic differences in how aggressively or conservatively agents consume an override budget.

C Model construction and derivations from Section 2

This appendix provides the complete construction of the gain from discretion Δ_t and the associated distributional definitions summarized in Section 2.

C.1 Allocation setup

In each period $t \in \{1, \dots, T\}$, two households arrive. Each household $i \in \{1, 2\}$ has a **baseline welfare** $w_i(t)$, drawn i.i.d. from a distribution $\mathcal{D}_w$, and a potential **welfare improvement** $I_i(t)$, drawn i.i.d. from a distribution I_θ, representing the gain in welfare if that household receives the resource. Baselines and improvements are mutually independent. The final welfare of a household that receives the resource is $w_i'(t) = w_i(t) + I_i(t)$. For any period, let m and M denote the indices of the households with minimum and maximum baseline welfare:

$$m = \arg\min_{i \in \{1,2\}} \{w_i(t)\}, \qquad M = \arg\max_{i \in \{1,2\}} \{w_i(t)\}.$$

C.2 Default policy and welfare-maximizing allocation

The *default policy* is a strictly needs-based rule: it always assigns the resource to the most vulnerable household (m) i.e., the household with the lower beginning welfare. This reflects standard prioritization logic in services such as homelessness or organ allocation [16, 43]. The total welfare under this policy is:

$$\mathrm{U}^{\text{pol}} = (w_m + I_m) + w_M = w_m + w_M + I_m.$$

In contrast, a welfare-maximizing assignment gives the resource to whichever household maximizes the total final welfare. For example, in homelessness services, while the administrative rule is to prioritize the most vulnerable, evaluation is often in terms of success in keeping households stably housed [23], which corresponds to prioritizing the household with the higher improvement. The maximum possible welfare under this is:

$$\mathrm{U}^{\text{max}} = w_m + w_M + \max\{I_m, I_M\}.$$

C.3 Derivation of the gain from discretion

The realized **gain from discretion** is the welfare difference between the two allocations:

$$\Delta = \mathrm{U}^{\text{max}} - \mathrm{U}^{\text{pol}} = \big(\max\{I_m, I_M\} - I_m\big) = (I_M - I_m)^+.$$

Note that the baseline welfare terms (w_m, w_M) cancel exactly, so the gain depends only on the improvements. The gain is non-negative by construction: $\Delta > 0$ if and only if $I_M > I_m$, i.e., when the less vulnerable household would benefit more from the resource than the neediest one. We call this event *misalignment*, denoted A_t^c, with complement A_t (alignment). Under i.i.d. continuous improvements, $\Pr(\mathsf{A}_t^c) = \Pr(\mathsf{A}_t) = 1/2$ by symmetry. We write $p = \Pr(\mathsf{A}_t)$.

C.4 Distributional definitions

The gain from discretion follows a mixture distribution across periods. In any period t:

$$\Delta_t = \begin{cases} 0, & \text{with probability } p \quad (\text{alignment} : I_m \geq I_M), \\ \overset{\text{iid}}{\sim} f_\Delta \text{ on } (0,\infty), & \text{with probability } 1-p \quad (\text{misalignment} : I_M > I_m), \end{cases}$$

where f_Δ is the density of the gain conditional on misalignment, with conditional CDF $F_{\Delta|A_t^c}$. We write F_Δ for the unconditional CDF, so that

$$F_\Delta(x) = p + (1-p)F_{\Delta|A_t^c}(x), \qquad x \geq 0.$$

Let $\mathrm{U}_t^{\text{pol}}$ and $\mathrm{U}_t^{\text{max}}$ respectively denote the period-t welfare under the policy-compliant and the welfare-maximizing allocation. By construction:

$$\mathrm{U}_t^{\text{max}} = \mathrm{U}_t^{\text{pol}} + \Delta_t.$$

Taking expectations and defining the per-period expected utilities:

$$\mu_{\text{pol}} = \mathbb{E}[\mathrm{U}_t^{\text{pol}}], \qquad \mu_{\text{max}} = \mathbb{E}[\mathrm{U}_t^{\text{max}}] = \mu_{\text{pol}} + \mathbb{E}[\Delta_t].$$

Because $\Delta_t = 0$ with probability p and is strictly positive with probability $1-p$, the unconditional expected gain decomposes as (mentioned in Section 2):

$$\mathbb{E}[\Delta_t] = (1-p)\mathbb{E}[\Delta_t \mid \mathsf{A}_t^c].$$

C.5 Simplification of the recursion from Section 2.2

We start from the conditional recursion, splitting on whether we have alignment (A_t with probability p) or not (A_t^c with probability $(1-p)$).

$$V(\tau,k) = \underbrace{p \cdot \mathbb{E}[(\mathrm{U}^{\text{max}} + V(\tau-1,k)) \mid \mathsf{A}_t]}_{\text{alignment } (\Delta=0)} + \underbrace{(1-p) \cdot \mathbb{E}\big[\max\big\{\big(\mathrm{U}^{\text{pol}} + V(\tau-1,k)\big), (\mathrm{U}^{\text{max}} + V(\tau-1,k-1))\big\} \mid \mathsf{A}_t^c\big]}_{\text{misalignment } (\Delta>0)}$$

Let us analyze the two portions in more detail:

- **On Alignment,** $\mathrm{U}^{\text{max}} = \mathrm{U}^{\text{pol}}$, so we can re-write the portion conditioned on A_t as:
$$\mathbb{E}[(\mathrm{U}^{\text{max}} + V(\tau-1,k)) \mid \mathsf{A}_t] = \mathbb{E}\big[\mathrm{U}^{\text{pol}} \mid \mathsf{A}_t\big] + V(\tau-1,k)$$
- **On Misalignment,** on the other hand, we can write $\mathrm{U}^{\text{max}} = \mathrm{U}^{\text{pol}} + \Delta$ with $\Delta > 0$. Then, looking at the $\max$ expression within the second half, we can re-write it as:
$$\begin{aligned} &\max\big\{\mathrm{U}^{\text{pol}} + V(\tau-1,k), \mathrm{U}^{\text{max}} + V(\tau-1,k-1)\big\} \\ &\quad = \mathrm{U}^{\text{pol}} + \max\{V(\tau-1,k), \Delta + V(\tau-1,k-1)\} \end{aligned}$$
Taking the conditional on A_t^c then gives us the second half as:
$$\begin{aligned} &= \mathbb{E}\big[\mathrm{U}^{\text{pol}} + \max\{V(\tau-1,k), \Delta + V(\tau-1,k-1)\} \mid \mathsf{A}_t^c\big] \\ &= \mathbb{E}\big[\mathrm{U}^{\text{pol}} \mid \mathsf{A}_t^c\big] + \mathbb{E}_{\Delta|\mathsf{A}_t^c}\left[\max\{V(\tau-1,k), \Delta + V(\tau-1,k-1)\}\right] \end{aligned}$$

Note that, by law of total expectation:

$$p \cdot \mathbb{E}\big[\mathrm{U}^{\text{pol}} \mid \mathsf{A}_t\big] + (1-p) \cdot \mathbb{E}\big[\mathrm{U}^{\text{pol}} \mid \mathsf{A}_t^c\big] = \mathbb{E}\big[\mathrm{U}^{\text{pol}}\big]. \tag{9}$$

Going back to the original recursion:

$$\begin{aligned} V(\tau,k) &= p \cdot \mathbb{E}[\mathrm{U}^{\max} + V(\tau-1,k) \mid \mathsf{A}_t] + \\ &\quad (1-p) \cdot \mathbb{E}\Big[\max\Big\{\Big(\mathrm{U}^{\text{pol}} + V(\tau-1,k)\Big), (\mathrm{U}^{\max} + V(\tau-1,k-1))\Big\} \mid \mathsf{A}_t^c\Big] \\ &= p \cdot \mathbb{E}\Big[\mathrm{U}^{\text{pol}} \mid \mathsf{A}_t\Big] + p \cdot V(\tau-1,k) + \\ &\quad (1-p) \cdot \mathbb{E}\Big[\mathrm{U}^{\text{pol}} \mid \mathsf{A}_t^c\Big] + (1-p) \cdot \mathbb{E}_{\Delta \mid \mathsf{A}_t^c}\left[\max\left\{V(\tau-1,k), \Delta + V(\tau-1,k-1)\right\}\right] \end{aligned}$$

Now using the result from Equation (9), we get our final recursion as:

$$V(\tau,k) = \mu_{\text{pol}} + pV(\tau-1,k) + (1-p)\mathbb{E}_{\Delta \mid \mathsf{A}_t^c}\Big[\max\big\{V(\tau-1,k), \Delta + V(\tau-1,k-1)\big\}\Big]. \tag{3}$$

D Proofs

D.1 Proof of Proposition 3.1

Proposition 3.1. *$c(\mathcal{D})$ admits a convenient closed form. Suppose $X \sim \mathcal{D}$ has non-negative support and finite mean. Then the base threshold is:*

$$c(\mathcal{D}) = \int_0^\infty F_{\mathcal{D}}(x)\big(1 - F_{\mathcal{D}}(x)\big)dx. \tag{6}$$

Proof. We can derive a compact expression for the base threshold $c(\mathcal{D})$. Let $X, X' \overset{\text{iid}}{\sim} \mathcal{D}$. By the law of total expectation, $c(\mathcal{D}) = \mathbb{E}[(X'-X)^+] = \mathbb{E}[\mathbb{E}[(X'-x)^+ \mid X = x]]$. For a fixed x, the inner conditional expectation is:

$$\mathbb{E}\big[(X'-x)^+ \mid X = x\big] = \int_x^\infty (y-x) f_{\mathcal{D}}(y) dy.$$

We use integration by parts. Let $g(y) = 1 - F_{\mathcal{D}}(y)$, so $g'(y) = -f_{\mathcal{D}}(y)$. Then:

$$\begin{aligned} \int_x^\infty (y-x) f_{\mathcal{D}}(y) dy &= -\int_x^\infty (y-x) g'(y) dy \\ &= -\Big[(y-x)g(y)\Big]_x^\infty + \int_x^\infty g(y) dy \\ &= 0 + \int_x^\infty (1 - F_{\mathcal{D}}(y)) dy. \qquad \text{(since } \lim_{y\to\infty} y(1-F_{\mathcal{D}}(y)) = 0\text{)} \end{aligned}$$

This equality holds assuming the mean is finite. Now, we take the outer expectation over $X = x$:

$$c(\mathcal{D}) = \mathbb{E}\left[\int_x^\infty (1 - F_{\mathcal{D}}(y)) dy\right] = \int_{-\infty}^\infty f_{\mathcal{D}}(x) \left(\int_x^\infty (1 - F_{\mathcal{D}}(y)) dy\right) dx.$$

The domain of integration is $\{(x,y) : -\infty < x \le y < \infty\}$. We swap the order of integration:

$$\begin{aligned} c(\mathcal{D}) &= \int_{-\infty}^\infty (1 - F_{\mathcal{D}}(y)) \left(\int_{-\infty}^y f_{\mathcal{D}}(x) dx\right) dy \\ &= \int_{-\infty}^\infty (1 - F_{\mathcal{D}}(y)) \cdot F_{\mathcal{D}}(y) dy. \end{aligned}$$

This gives the compact expression for c in terms of the CDF alone:

$$c(\mathcal{D}) = \int_{-\infty}^\infty F_{\mathcal{D}}(x)(1 - F_{\mathcal{D}}(x)) dx.$$

Because X has non-negative support, the integral runs from 0 to ∞.

$$c(\mathcal{D}) = \int_0^\infty F_{\mathcal{D}}(x)(1 - F_{\mathcal{D}}(x)) dx.$$

□

D.2 Proof of Theorem 3.1 (general behavioral invariance)

This section provides the complete proof of Theorem 3.1. We first derive the W-recursion, then prove each part.

Theorem 3.1 (General behavioral invariance)**.** *Let per-household improvements follow a location-scale family $I = a + sX$ with $s > 0$, where the base shape $\mathcal{D}$ has non-negative support and finite mean. Then for every state (τ, k) with $\tau \geq 1$ and $k \geq 1$:*

(i) ***The discretionary value scales linearly with s and is invariant to a:*** *there exists a quantity $w(\tau, k; \mathcal{D})$, depending only on $\mathcal{D}$, p, and (τ, k), such that $W(\tau, k) = s \cdot w(\tau, k; \mathcal{D})$.*

(ii) ***The optimal threshold scales linearly with s and is invariant to a:*** *$\mathcal{T}_{\tau,k} = s \cdot \theta_{\tau,k}(\mathcal{D})$, where $\theta_{\tau,k}(\mathcal{D}) = w(\tau - 1, k; \mathcal{D}) - w(\tau - 1, k - 1; \mathcal{D})$ is the base threshold.*

(iii) ***The spending probability at every state is scale- and location-free:*** *$q_{\tau,k} = 1 - G_{\mathcal{D}}(\theta_{\tau,k}(\mathcal{D}))$, where $G_{\mathcal{D}}$ is the CDF of the base difference $\Delta_0 = X' - X$ as defined in* (5)*.*

The proof proceeds by induction on τ through the recursion as detailed in Equation (3). For this argument to work, we need two ingredients:

1. A recursion for the "discretionary value" – the part of the value function that comes from having override budget, stripped of the baseline utility that accrues regardless. This recursion is where the induction operates.
2. The factorization $\Delta = s \cdot Z$, which shows that the gain from discretion separates into a scale factor s and a shape-only random variable Z. This is the engine of invariance: it is what allows s to be pulled out of the recursion at every step.

D.2.1 Derivation of the W-recursion from Equation (7) in Section 3.2

The recursion in Equation (3) mixes two sources of value: the baseline utility μ_{pol} that is accrued every period regardless of whether discretion is exercised, and the incremental value that arises from having override budget available. To isolate the latter, we subtract out the baseline accumulation.

Definition 1 (Discretionary value)**.** *The* discretionary value *at state (τ, k) is $W(\tau, k) = V(\tau, k) - \tau \cdot \mu_{pol}$. This measures the expected welfare gain from holding k override units with τ periods remaining, relative to the baseline of following the default policy in every remaining period.*

With this definition, we state and prove the following lemma that gives us Equation (7).

Lemma D.1 (Recursion for W)**.** *The discretionary value satisfies*

$$W(\tau, k) = p \cdot W(\tau - 1, k) + (1 - p) \cdot \mathbb{E}\left[\max\left\{W(\tau - 1, k), \Delta + W(\tau - 1, k - 1)\right\} \middle| \mathsf{A}_t^c\right] \quad (7)$$

with boundary conditions $W(0, k) = 0$ for all k and $W(\tau, 0) = 0$ for all $\tau \geq 0$.

Proof. We substitute $V(\tau, k) = W(\tau, k) + \tau \cdot \mu_{\text{pol}}$ into every occurrence of V in the recursion from Equation (3). The left-hand side becomes $W(\tau, k) + \tau \cdot \mu_{\text{pol}}$. On the right-hand side, every $V(\tau - 1, \cdot)$ term gains an additive offset of $(\tau - 1) \cdot \mu_{\text{pol}}$. Inside the $\max$ operator, this offset appears in both arguments and therefore factors out. Collecting the constant terms on the right gives $\mu_{\text{pol}} + p \cdot (\tau - 1)\mu_{\text{pol}} + (1 - p) \cdot (\tau - 1)\mu_{\text{pol}} = \tau \cdot \mu_{\text{pol}}$. Canceling $\tau \cdot \mu_{\text{pol}}$ from both sides yields (7).

For the boundary conditions: $V(0, k) = 0$ gives $W(0, k) = 0$. When the budget is exhausted, $V(\tau, 0) = \tau \cdot \mu_{\text{pol}}$, giving $W(\tau, 0) = 0$. □

The recursion (7) further admits a useful simplification. Write $d_k = W(\tau - 1, k) - W(\tau - 1, k - 1)$ for the marginal value of the k-th budget unit at time $\tau - 1$. Since $W(\tau - 1, k) = W(\tau - 1, k - 1) + d_k$, the $\max$ inside (7) becomes $W(\tau - 1, k - 1) + \max\{d_k, \Delta\}$. Substituting and simplifying:

$$W(\tau, k) = W(\tau - 1, k) + (1 - p) \cdot \mathbb{E}\left[(\Delta - d_k)^+ \middle| \mathsf{A}_t^c\right]. \quad (10)$$

The interpretation is direct: the discretionary value at (τ, k) equals the value at $(\tau - 1, k)$ – carrying the same budget into one fewer period – plus the expected "bonus" from the current period, which is collected only when a misalignment occurs and the gain exceeds the opportunity cost of spending.

D.2.2 The gain factorization from Section 3.2

The second ingredient is a factorization of the gain Δ under the location-scale assumption. Recall that the gain from discretion is $\Delta = (I_M - I_m)^+$, where $I_M = a + sX_M$ and $I_m = a + sX_m$ are the improvement draws for the two households, with $X_M, X_m \overset{\text{iid}}{\sim} \mathcal{D}$. Therefore:

$$\Delta = ((a + sX_M) - (a + sX_m))^+ = (s(X_M - X_m))^+ = s \cdot (X_M - X_m)^+.$$

Two things happen. First, the location parameter a cancels: it appears with the same sign in both improvements, so it drops out of the difference. Second, the scale parameter s factors out of the positive-part operator because $s > 0$ preserves sign. Writing $Z = (X_M - X_m)^+$, whose distribution depends only on the base shape $\mathcal{D}$, we have

$$\Delta = s \cdot Z$$

This is the fundamental factorization: the gain in any period is the scale s times a shape-only random variable Z.

D.2.3 Putting the pieces together for the proof

Now we are ready to prove the Theorem 3.1. The argument for part (i) is the core of the proof and proceeds by induction on the number of periods remaining τ. It combines the ingredients developed above: the W-recursion (Lemma D.1), the gain factorization, and the ability to factor s out of the resulting expression. Parts (ii) and (iii) then follow as short algebraic consequences.

Proof of Part (i): $W(\tau, k) = s \cdot w(\tau, k; \mathcal{D})$. We prove by induction on τ that, for all $k \in \{0, 1, \ldots, K\}$ (for $k = 0$ this is immediate from $W(\tau, 0) = 0$, so the inductive step below takes $k \geq 1$), the discretionary value $W(\tau, k)$ can be written as s times a quantity $w(\tau, k; \mathcal{D})$ that does not depend on a or s.

Step 1: Base case ($\tau = 0$). By the boundary condition in Lemma D.1, $W(0, k) = 0$ for all k. We can write this as $W(0, k) = s \cdot 0$, so taking $w(0, k; \mathcal{D}) = 0$ satisfies the claim trivially: both sides are zero regardless of s and a.

Step 2: Inductive hypothesis. Assume the claim holds at time $\tau - 1$. That is, assume there exist quantities $w(\tau - 1, j; \mathcal{D})$, depending only on $\mathcal{D}$, p, and j, such that

$$W(\tau - 1, j) = s \cdot w(\tau - 1, j; \mathcal{D}) \qquad \text{for all } j \in \{0, 1, \ldots, K\}. \tag{11}$$

Step 3: Substituting into the W-recursion. Our goal is to show that $W(\tau, k)$ also has this form. We start from the recursion established in Lemma D.1:

$$W(\tau, k) = p \cdot W(\tau - 1, k) + (1 - p) \cdot \mathbb{E}_{\Delta | \mathsf{A}_t^c} \left[\max \left\{ W(\tau - 1, k), \Delta + W(\tau - 1, k - 1) \right\} \right].$$

We now replace every quantity on the right-hand side using the inductive hypothesis (11) for the W terms, and the gain factorization $\Delta = s \cdot Z$ for the gain. Making both substitutions simultaneously, we get:

$$\begin{aligned} W(\tau, k) = {} & p \cdot s \cdot w(\tau - 1, k; \mathcal{D}) \\ & + (1 - p) \cdot \mathbb{E}_{\Delta | \mathsf{A}_t^c} \left[\max \left\{ s \cdot w(\tau - 1, k; \mathcal{D}), s \cdot Z + s \cdot w(\tau - 1, k - 1; \mathcal{D}) \right\} \right]. \end{aligned} \tag{12}$$

At this stage, s appears in every term inside the $\max$ as well as in the first summand. The rest of the argument consists of extracting it.

Step 4: Factoring s out of the $\max$. Look at the two arguments inside the $\max$ in (12). The first argument is $s \cdot w(\tau - 1, k; \mathcal{D})$. The second argument is $s \cdot Z + s \cdot w(\tau - 1, k - 1; \mathcal{D}) = s \cdot [Z + w(\tau - 1, k - 1; \mathcal{D})]$. Both arguments have s as a common factor. Because $s > 0$, multiplying by s preserves the ordering of real numbers, so the identity $\max\{s \cdot u, s \cdot v\} = s \cdot \max\{u, v\}$ holds. Applying this:

$$\begin{aligned} & \max \left\{ s \cdot w(\tau - 1, k; \mathcal{D}), s \cdot Z + s \cdot w(\tau - 1, k - 1; \mathcal{D}) \right\} \\ & \quad = s \cdot \max \left\{ w(\tau - 1, k; \mathcal{D}), Z + w(\tau - 1, k - 1; \mathcal{D}) \right\}. \end{aligned} \tag{13}$$

Step 5: Pulling s outside the expectation. Substituting (13) back into (12), the expectation becomes

$$\mathbb{E}_{\Delta|\mathsf{A}_t^c}\left[s \cdot \max\left\{w(\tau-1,k;\mathcal{D}), Z + w(\tau-1,k-1;\mathcal{D})\right\}\right].$$

Since s is a positive constant (it does not depend on the random variable Z), the linearity of expectation allows us to pull it outside:

$$= s \cdot \mathbb{E}_{\Delta|\mathsf{A}_t^c}\left[\max\left\{w(\tau-1,k;\mathcal{D}), Z + w(\tau-1,k-1;\mathcal{D})\right\}\right].$$

Substituting this back, the full expression (12) becomes:

$$\begin{aligned} W(\tau,k) = p \cdot s \cdot w(\tau-1,k;\mathcal{D}) \\ + (1-p) \cdot s \cdot \mathbb{E}_{\Delta|\mathsf{A}_t^c}\left[\max\left\{w(\tau-1,k;\mathcal{D}), Z + w(\tau-1,k-1;\mathcal{D})\right\}\right]. \end{aligned} \tag{14}$$

Step 6: Factoring s from the entire expression. Both terms on the right-hand side of (14) have s as a common factor. Extracting it:

$$\begin{aligned} W(\tau,k) &= s \cdot w(\tau,k;\mathcal{D}), \qquad \text{where} \\ w(\tau,k;\mathcal{D}) &= p \cdot w(\tau-1,k;\mathcal{D}) \\ &\quad + (1-p) \cdot \mathbb{E}_{\Delta|\mathsf{A}_t^c}\left[\max\left\{w(\tau-1,k;\mathcal{D}), Z + w(\tau-1,k-1;\mathcal{D})\right\}\right]. \end{aligned} \tag{15}$$

The last two lines define $w(\tau,k;\mathcal{D})$.

Step 7: Verifying that $w(\tau,k;\mathcal{D})$ is free of a and s. For the induction to close, we must confirm that $w(\tau,k;\mathcal{D})$ as defined in (15) depends only on the base shape $\mathcal{D}$, the alignment probability p, and the state (τ,k). We check each component:

- The quantities $w(\tau-1,k;\mathcal{D})$ and $w(\tau-1,k-1;\mathcal{D})$ are free of a and s by the inductive hypothesis.
- The random variable $Z = (X_M - X_m)^+$ is determined entirely by two independent draws from the base shape $\mathcal{D}$. It does not involve a (which canceled in the difference) or s (which was factored out).
- The alignment probability p is a model primitive.
- No other quantities appear.

Therefore $w(\tau,k;\mathcal{D})$ is indeed free of a and s. This completes the inductive step, and by induction, part (i) holds for all $\tau \in \{0,1,\ldots,T\}$ and all $k \in \{0,1,\ldots,K\}$. □

Proof of Part (ii): $\mathcal{T}_{\tau,k} = s \cdot \theta_{\tau,k}(\mathcal{D})$. This part follows directly from part (i) by taking differences. The optimal threshold is defined as the marginal value of carrying one additional unit of budget into the future:

$$\mathcal{T}_{\tau,k} = V(\tau-1,k) - V(\tau-1,k-1).$$

Since $V(\tau-1,j) = W(\tau-1,j) + (\tau-1)\mu_{\text{pol}}$ (by Definition 1), the baseline term $(\tau-1)\mu_{\text{pol}}$ appears in both $V(\tau-1,k)$ and $V(\tau-1,k-1)$ and cancels in the difference:

$$\mathcal{T}_{\tau,k} = W(\tau-1,k) - W(\tau-1,k-1).$$

Now apply part (i) to both W terms:

$$\begin{aligned} \mathcal{T}_{\tau,k} &= s \cdot w(\tau-1,k;\mathcal{D}) - s \cdot w(\tau-1,k-1;\mathcal{D}) \\ &= s \cdot \left[w(\tau-1,k;\mathcal{D}) - w(\tau-1,k-1;\mathcal{D})\right]. \end{aligned} \tag{16}$$

Define the *base threshold* $\theta_{\tau,k}(\mathcal{D}) = w(\tau-1,k;\mathcal{D}) - w(\tau-1,k-1;\mathcal{D})$. This is a difference of two shape-only quantities, so it inherits their invariance. We conclude $\mathcal{T}_{\tau,k} = s \cdot \theta_{\tau,k}(\mathcal{D})$. □

Proof of Part (iii): $q_{\tau,k} = 1 - G_{\mathcal{D}}\left(\theta_{\tau,k}(\mathcal{D})\right)$. This part shows that the scale s cancels from the spending probability. The argument has three sub-steps: we express the CDF of Δ in terms of the base-shape CDF, substitute the scaled threshold from part (ii), and observe that s cancels.

Step 1: Relating F_Δ to $G_\mathcal{D}$. Recall from Section D.2.2 that the unconditional gain is $\Delta = s \cdot Z'$, where $Z' = \Delta_0^+ = (X_M - X_m)^+$ is the positive part of the base difference. (Note that Z' takes value 0 under alignment and is strictly positive under misalignment, mirroring the mixture structure of Δ.) For any $x \geq 0$:

$$F_\Delta(x) = \Pr(\Delta \leq x) = \Pr(s \cdot Z' \leq x) = \Pr\left(Z' \leq \frac{x}{s}\right) = G_\mathcal{D}\left(\frac{x}{s}\right), \tag{17}$$

where $G_\mathcal{D}$ is the CDF of the signed difference $X_M - X_m$. For nonnegative arguments, this agrees with the CDF of its positive part Z'.

Step 2: Evaluating at the threshold. We substitute $x = \mathcal{T}_{\tau,k} = s \cdot \theta_{\tau,k}(\mathcal{D})$ from part (ii) into (17):

$$F_\Delta(\mathcal{T}_{\tau,k}) = G_\mathcal{D}\left(\frac{s \cdot \theta_{\tau,k}(\mathcal{D})}{s}\right) = G_\mathcal{D}\left(\theta_{\tau,k}(\mathcal{D})\right). \tag{18}$$

The scale parameter s appears in the numerator (from the threshold scaling in part (ii)) and in the denominator (from the CDF rescaling in (17)), and cancels exactly.

Step 3: Concluding. The spending probability is therefore

$$q_{\tau,k} = 1 - F_\Delta(\mathcal{T}_{\tau,k}) = 1 - G_\mathcal{D}\left(\theta_{\tau,k}(\mathcal{D})\right). \tag{19}$$

The right-hand side involves only the base-shape CDF $G_\mathcal{D}$ evaluated at the base threshold $\theta_{\tau,k}(\mathcal{D})$. Both of these are functions of $\mathcal{D}$, p, and (τ, k) alone. The location parameter a never entered the gain Δ (it canceled in the difference in Section D.2.2), and the scale parameter s has canceled in (18). □

This completes the proof of Theorem 3.1. □

E Two-period special case $(T = 2, K = 1)$

This appendix works out the budgeted discretion problem in the minimal non-trivial case – a single override ($K = 1$) over two periods ($T = 2$) – as defined in Section 3.3, providing closed-form expressions that are recovered as special cases of the general results in Section 3.

E.1 State space and terminal values

The state is (τ, k) with $\tau \in \{1, 2\}$ periods remaining and $k \in \{0, 1\}$ budget units. The four states and their values are:

τ	k	$V(\tau, k)$	
1	0	μ_{pol}	(no budget: collect default welfare)
1	1	μ_{max}	(budget available: override any misalignment)
2	0	$2\mu_{\text{pol}}$	(no budget for either period)
2	1	solved below	

The non-trivial case is $(\tau, k) = (2, 1)$: the agent has two periods remaining and one override to allocate.

E.2 Derivation of the optimal threshold

Applying the recursion in Equation (3) at state $(2, 1)$:

$$V(2,1) = \mu_{\text{pol}} + pV(1,1) + (1-p)\mathbb{E}_{\Delta|\mathsf{A}_t^c}\left[\max\left\{V(1,1), \Delta + V(1,0)\right\}\right].$$

Substituting $V(1,1) = \mu_{\text{max}}$ and $V(1,0) = \mu_{\text{pol}}$:

$$V(2,1) = \mu_{\text{pol}} + p\mu_{\text{max}} + (1-p)\mathbb{E}\left[\max\left\{\mu_{\text{max}}, \Delta + \mu_{\text{pol}}\right\} \middle| \mathsf{A}_t^c\right].$$

The agent overrides in the first period if and only if the immediate gain exceeds the value of saving the budget for the second period. By the general threshold rule (4), this requires $\Delta_1 \geq \mathcal{T}_{2,1}$, where:

$$\mathcal{T}_{2,1} = V(1,1) - V(1,0) = \mu_{\max} - \mu_{\text{pol}} = \mathbb{E}[\Delta].$$

Since $\mathbb{E}[\Delta] = (1-p)\mathbb{E}[\Delta \mid \mathsf{A}_t^c]$, the threshold decomposes as:

$$\mathcal{T}_{2,1} = (1-p) \cdot \mathbb{E}[\Delta \mid \mathsf{A}_t^c].$$

The threshold is the product of the *frequency* of override opportunities $(1-p)$ and their expected *magnitude* ($\mathbb{E}[\Delta \mid \mathsf{A}_t^c]$). Intuitively, if misalignments are rare (p large), the threshold is low because the option value of holding the budget is small; if misalignments are frequent and large, the threshold is high because saving the override for the second period is more valuable.

In the final period ($\tau = 1$), if budget remains the agent overrides whenever a misalignment occurs (i.e., whenever $\Delta_2 > 0$), since there is no future period to save for.

E.3 The patience scalar ψ

While the threshold $\mathcal{T}_{2,1}$ depends on the units of welfare, the agent's observable behavior–how often they override–is unit-free. Let ψ denote the probability that the agent does *not* spend the budget in the first period. The agent saves in two scenarios:

1. **Alignment** (probability p): the default policy is already optimal ($\Delta_1 = 0$), so there is no gain to capture.
2. **Insufficient misalignment** (probability $(1-p) \cdot F_{\Delta|\mathsf{A}_t^c}(\mathcal{T}_{2,1})$): a misalignment occurs but the realized gain falls below the threshold.

Combining these disjoint events:

$$\psi = p + (1-p)F_{\Delta|\mathsf{A}_t^c}(\mathcal{T}_{2,1}).$$

By the law of total probability, this is exactly the unconditional CDF of Δ evaluated at the threshold:

$$\psi = F_\Delta(\mathcal{T}_{2,1}) = \Pr(\Delta_1 \leq \mathcal{T}_{2,1}).$$

The probability that the optimal agent exercises discretion in the first period is $1 - \psi$.

E.4 Conditional spending probabilities

We first establish the conditional spending probabilities used in the behavioral comparisons of Appendix E.8.

- **Oracle Policy ($\tilde{\pi}$):** An oracle sees both Δ_1 and Δ_2 in advance. It spends at time t if either (i) only period t misaligns, or (ii) both periods misalign and $\Delta_t > \Delta_s$ (where $s = 3 - t$). Conditional on both periods misaligning, the positive gains are independent and identically distributed with a continuous distribution. Hence $\Pr(\Delta_t > \Delta_s \mid A_t^c, A_s^c) = 1/2$.

$$\overset{\tilde{\pi}}{\Pr}(\text{spend at } t \mid \mathsf{A}_t^c) = p \cdot 1 + (1-p) \cdot \frac{1}{2} = \frac{1+p}{2}. \tag{20}$$

- **DP Policy (π^*):** The DP only sees the current state.
  - At $t = 1$, it spends if Δ_1 clears the threshold:

$$\overset{\pi^*}{\Pr}(\text{spend at } 1 \mid \mathsf{A}_1^c) = \Pr(\Delta_1 \geq \mathcal{T}_{2,1} \mid \mathsf{A}_1^c) = 1 - F_{\Delta|\mathsf{A}_1^c}(\mathcal{T}_{2,1}) = \frac{1-\psi}{1-p}. \tag{21}$$

  - At $t = 2$, it spends upon any misalignment *if the budget is available*. The budget is available if $t = 1$ was an alignment (prob. p) or if it was a misalignment but the gain was too low to spend (probability $(1-p)F_{\Delta|\mathsf{A}_1^c}(\mathcal{T}_{2,1})$)

$$\overset{\pi^*}{\Pr}(\text{spend at } 2 \mid \mathsf{A}_2^c) = p + (1-p)F_{\Delta|\mathsf{A}_1^c}(\mathcal{T}_{2,1}) = \psi. \tag{22}$$

E.5 Proof of Theorem 3.2 (two-period invariance)

Theorem 3.2 (Shape dependence of optimal policy, two-period)**.** *Let the improvement I follow a location–scale family $I = a + sX$ with $s > 0$, where the base shape $\mathcal{D}$ has non-negative support and finite mean. Then, when $T = 2$ and $K = 1$:*

(i) ***Optimal threshold scales linearly with*** *s. The two-period threshold $\mathcal{T}_{2,1} = \mathbb{E}[\Delta]$ is invariant to location a and satisfies $\mathcal{T}_{2,1} = s \cdot c(\mathcal{D})$, where the base threshold $c(\mathcal{D})$ depends only on the shape and is given by* (6)*.*

(ii) ***Policy behavior is summarized by a single scalar.*** *All optimal spending probabilities are functions of $\psi(\mathcal{D}) = F_\Delta(\mathcal{T}_{2,1})$, which is invariant to a and s and equals $\psi(\mathcal{D}) = G_\mathcal{D}(c(\mathcal{D}))$.*

Proof. We prove the theorem by showing that for any continuous base distribution $\mathcal{D}$ with finite mean, and for any location-scale transform defined by parameters (a, s) with $s > 0$, the key probability ψ is invariant. Let I_1, I_2 be two independent draws from the location-scale family $I = a + sX$, where $X_1, X_2 \overset{\text{iid}}{\sim} \mathcal{D}$ are drawn from the base shape. We proceed in four steps.

Step 1: Differences cancel location and scale linearly. The signed difference between two realizations under the location-scale transform is:

$$I_2 - I_1 = (a + sX_2) - (a + sX_1) = s(X_2 - X_1) = s\Delta_0.$$

The location parameter a cancels out entirely, and the scale s factors out. Consequently, the realized gain $\Delta = (I_M - I_m)^+$ scales as $\Delta = s\Delta_0^+$.

Step 2: The optimal threshold scales by s. The optimal threshold $\mathcal{T}_{2,1}$ is defined as the expectation of the realized gain Δ. Since $s > 0$:

$$\mathcal{T}_{2,1} = \mathbb{E}[\Delta] = \mathbb{E}\big[s\Delta_0^+\big] = s \cdot \mathbb{E}\big[\Delta_0^+\big] = s \cdot c(\mathcal{D}).$$

This establishes Part 1 of the theorem: the threshold is the base threshold $c(\mathcal{D})$ scaled linearly by s.

Step 3: The CDF of the difference rescales its argument. Let G_I be the CDF of the difference $I_2 - I_1$. Using the result from Step 1:

$$G_I(x) = \Pr(I_2 - I_1 \le x) = \Pr(s\Delta_0 \le x) = \Pr\left(\Delta_0 \le \frac{x}{s}\right) = G_\mathcal{D}\Big(\frac{x}{s}\Big).$$

Step 4: Evaluate at the threshold to obtain invariance. We now combine these results to find the patience scalar ψ, the probability of not spending in period 1. Since $F_\Delta(x) = G_I(x)$ for $x \ge 0$ (the pair (I_M, I_m) has the same distribution as (I_2, I_1), and $\{(I_M - I_m)^+ \le x\} = \{I_M - I_m \le x\}$ for $x \ge 0$):

$$\begin{aligned}
\psi &= F_\Delta(\mathcal{T}_{2,1}) = G_I(\mathcal{T}_{2,1}) && \text{(definition of } \psi\text{)} \\
&= G_I\big(s \cdot c(\mathcal{D})\big) && \text{(from Step 2)} \\
&= G_\mathcal{D}\Big(\frac{s \cdot c(\mathcal{D})}{s}\Big) && \text{(from Step 3)} \\
&= G_\mathcal{D}\big(c(\mathcal{D})\big).
\end{aligned}$$

This establishes Part 2 of the theorem: ψ depends strictly on the base shape objects $G_\mathcal{D}$ and $c(\mathcal{D})$, and is invariant to the location a and scale s.

□

E.6 Recovery of $(T = 2, K = 1)$ invariance from the general Theorem 3.1

As mentioned at the end of Section 3.3, Theorem 3.2 can be proved either directly, as in Appendix E.5, or as a corollary of the more general Theorem 3.1. Under the location–scale model $I = a + sX$ with $X \sim \mathcal{D}$, the general theorem gives $W(\tau, k) = s \cdot w(\tau, k; \mathcal{D})$. At $\tau = 1$:

$$\begin{aligned}
&W(1, 0) = 0, \\
&W(1, 1) = (1 - p)\mathbb{E}[\Delta \mid \mathsf{A}_t^c] = s(1 - p)\mathbb{E}[Z \mid Z > 0],
\end{aligned}$$

where $Z = (X' - X)^+$ for $X, X' \overset{\text{iid}}{\sim} \mathcal{D}$. Since $Z = 0$ with probability p and $Z > 0$ with probability $1 - p$, the unconditional expectation satisfies $\mathbb{E}[Z] = (1 - p)\mathbb{E}[Z \mid Z > 0]$. Therefore:

$$w(1, 1; \mathcal{D}) = (1 - p)\mathbb{E}[Z \mid Z > 0] = \mathbb{E}[Z] = \mathbb{E}[(X' - X)^+] = c(\mathcal{D}).$$

The base threshold is $\theta_{2,1}(\mathcal{D}) = w(1, 1; \mathcal{D}) - w(1, 0; \mathcal{D}) = c(\mathcal{D}) - 0 = c(\mathcal{D})$. By Theorem 3.1(iii), the spending probability is:

$$q_{2,1} = 1 - G_{\mathcal{D}}(c(\mathcal{D})) = 1 - \psi,$$

confirming that $\psi = G_{\mathcal{D}}(c(\mathcal{D}))$ depends only on the base shape $\mathcal{D}$ and is invariant to both a and s.

E.7 Verification and behavioral examples

This appendix verifies the invariance and shape-dependence results from Section 3 through analytical closed forms and numerical experiments.

E.7.1 Analytical example: Exponential improvements.

Let the base improvements follow an Exponential distribution, $X \sim \text{Exp}(\lambda)$. This is a pure scale family where the scale parameter is $b = 1/\lambda$. We invoke the standard result that the difference of two i.i.d. exponential random variables follows a Laplace distribution with location 0 and scale b. Thus, the base difference density is $f_{\Delta_0}(x) = \frac{1}{2b} e^{-|x|/b}$. Using this explicit form, we calculate the base threshold c:

$$c = \mathbb{E}[\Delta_0^+] = \int_0^\infty x \cdot \frac{1}{2b} e^{-x/b} dx = \frac{b}{2}.$$

We then evaluate the Laplace CDF, $F(x) = 1 - \frac{1}{2} e^{-x/b}$ (for $x \geq 0$), at this threshold:

$$\psi_{\text{Exp}} = F_{\text{Laplace}}(b/2) = 1 - \frac{1}{2} e^{-(b/2)/b} = 1 - \frac{1}{2} e^{-1/2} \approx 0.697.$$

E.7.2 Analytical example: Half-Normal improvements.

Let $X \sim \text{HalfNormal}(\sigma)$. The difference of two Half-Normals does not follow a named textbook distribution, but its CDF can be derived via convolution. For $X_1, X_2 \overset{\text{iid}}{\sim} \text{HalfNormal}(\sigma)$, we compute the base threshold $c(\mathcal{D}) = \mathbb{E}[(\Delta_0)^+]$ in closed form. Let us first calculate the pdf of the difference between two half normals. The half-normal PDF and CDF are

$$f_X(x) = \begin{cases} \frac{\sqrt{2}}{\sigma\sqrt{\pi}} \exp\left(-\frac{x^2}{2\sigma^2}\right), & x \geq 0, \\ 0, & x < 0, \end{cases} \qquad F_{\mathcal{D}}(x) = \Pr(X \leq x) = \text{erf}\left(\frac{x}{\sqrt{2}\sigma}\right), \; x \geq 0,$$

where erf is the standard error function. The density of the difference $\Delta_0 = X_2 - X_1$ is given by the usual convolution:

$$f_{\Delta_0}(d) = \int_{-\infty}^{\infty} f_X(x) f_X(x + d)\, dx = \int_{\max\{0, -d\}}^{\infty} f_X(x) f_X(x + d)\, dx.$$

For $d \geq 0$ we have

$$\begin{aligned} f_{\Delta_0}(d) &= \int_0^\infty \frac{\sqrt{2}}{\sigma\sqrt{\pi}} \exp\left(-\frac{x^2}{2\sigma^2}\right) \cdot \frac{\sqrt{2}}{\sigma\sqrt{\pi}} \exp\left(-\frac{(x+d)^2}{2\sigma^2}\right) dx \\ &= \frac{2}{\pi\sigma^2} \int_0^\infty \exp\left(-\frac{x^2 + (x+d)^2}{2\sigma^2}\right) dx \\ &= \frac{2}{\pi\sigma^2} \exp\left(-\frac{d^2}{4\sigma^2}\right) \int_0^\infty \exp\left(-\frac{(x + d/2)^2}{\sigma^2}\right) dx. \end{aligned}$$

With the change of variables $u = (x + d/2)/\sigma$, $dx = \sigma\, du$ and lower limit $u = d/(2\sigma)$, this becomes

$$f_{\Delta_0}(d) = \frac{2}{\pi\sigma^2} \exp\left(-\frac{d^2}{4\sigma^2}\right) \sigma \int_{d/(2\sigma)}^\infty e^{-u^2}\, du.$$

Using the complementary error function

$$\operatorname{erfc}(z) = \frac{2}{\sqrt{\pi}} \int_z^\infty e^{-u^2}\, du,$$

we obtain, for $d \geq 0$,

$$f_{\Delta_0}(d) = \frac{1}{\sqrt{\pi}\sigma} \exp\left(-\frac{d^2}{4\sigma^2}\right) \operatorname{erfc}\left(\frac{d}{2\sigma}\right). \tag{23}$$

Since X_1 and X_2 are i.i.d., the distribution of Δ_0 is symmetric around zero, so $f_{\Delta_0}(-d) = f_{\Delta_0}(d)$ and $G_{\mathcal{D}}(0) = 1/2$. So, for $z \geq 0$, the CDF can be written as

$$G_{\mathcal{D}}(z) = \Pr(\Delta_0 \leq z) = \frac{1}{2} + \int_0^z f_{\Delta_0}(d)\, dd.$$

Next consider,

$$H(z) = 1 - \frac{1}{2} \operatorname{erfc}^2\left(\frac{z}{2\sigma}\right), \quad z \geq 0.$$

We can use the identity $\frac{d}{dz} \operatorname{erfc}(z) = \frac{-2}{\sqrt{\pi}} \exp(-z^2)$ and the chain rule to differentiate $H(z)$ with respect to z to get:

$$H'(z) = \frac{1}{\sqrt{\pi}\sigma} \exp\left(-\frac{z^2}{4\sigma^2}\right) \operatorname{erfc}\left(\frac{z}{2\sigma}\right) = f_{\Delta_0}(z), \quad z \geq 0,$$

and $H(0) = 1/2 = G_{\mathcal{D}}(0)$. Thus $G_{\mathcal{D}}(z) = H(z)$ for all $z \geq 0$, yielding

$$G_{\mathcal{D}}(z) = 1 - \frac{1}{2} \operatorname{erfc}^2\left(\frac{z}{2\sigma}\right) \quad z \geq 0.$$

By symmetry, $G_{\mathcal{D}}(-z) = 1 - G_{\mathcal{D}}(z)$ for $z > 0$.

Now, we derive a closed form for $c(\mathcal{D})$. Recall that

$$c(\mathcal{D}) = \mathbb{E}\left[(\Delta_0)^+\right] = \int_0^\infty d\, f_{\Delta_0}(d)\, dd.$$

Substituting (23) and changing variables $u = d/(2\sigma)$, so $d = 2\sigma u$ and $dd = 2\sigma\, du$, gives

$$\begin{aligned} c(\mathcal{D}) &= \int_0^\infty d \cdot \frac{1}{\sqrt{\pi}\sigma} \exp\left(-\frac{d^2}{4\sigma^2}\right) \operatorname{erfc}\left(\frac{d}{2\sigma}\right) dd \\ &= \frac{4\sigma}{\sqrt{\pi}} \int_0^\infty u e^{-u^2} \operatorname{erfc}(u)\, du. \end{aligned} \tag{24}$$

To evaluate the integral, write

$$\operatorname{erfc}(u) = \frac{2}{\sqrt{\pi}} \int_u^\infty e^{-t^2}\, dt,$$

so that

$$\begin{aligned} \int_0^\infty u e^{-u^2} \operatorname{erfc}(u)\, du &= \frac{2}{\sqrt{\pi}} \int_0^\infty \int_u^\infty u e^{-u^2} e^{-t^2}\, dt\, du \\ &= \frac{2}{\sqrt{\pi}} \int_0^\infty e^{-t^2} \left[\int_0^t u e^{-u^2}\, du\right] dt, \end{aligned}$$

where we have interchanged the order of integration. The inner integral is

$$\int_0^t u e^{-u^2}\, du = \frac{1}{2}\left(1 - e^{-t^2}\right),$$

so

$$\begin{aligned} \int_0^\infty u e^{-u^2} \operatorname{erfc}(u)\, du &= \frac{1}{\sqrt{\pi}} \int_0^\infty \left(e^{-t^2} - e^{-2t^2}\right) dt \\ &= \frac{1}{\sqrt{\pi}} \left[\frac{1}{2}\sqrt{\pi} - \frac{1}{2}\sqrt{\frac{\pi}{2}}\right] = \frac{2 - \sqrt{2}}{4}. \end{aligned}$$

Substituting back into (24) yields

$$c(\mathcal{D}) = \frac{4\sigma}{\sqrt{\pi}} \cdot \frac{2-\sqrt{2}}{4} = \frac{\sigma}{\sqrt{\pi}} \left(2-\sqrt{2}\right)$$

Substituting this into the derived CDF:

$$\psi_{\text{HN}} = 1 - \frac{1}{2}\,\text{erfc}^2\left(\frac{2-\sqrt{2}}{2\sqrt{\pi}}\right) \approx 0.668.$$

In both cases, the optimal spending probability is not dependent on the parameters of the distribution, confirming Theorem 3.2.

E.7.3 Numerical verification: scale vs. shape

This appendix reports the Monte Carlo simulations used to verify Theorem 3.2, which says that ψ should remain fixed if we stretch the distribution (scale change) but move if we alter its tail profile (shape change). All estimates use $N = 10^6$ independent draws per configuration.

Scale invariance. Table 3 empirically validates the scale-invariance prediction: changing the scale of the improvement distribution rescales the optimal cutoff $\hat{c}$ but leaves the unconditional patience scalar $\hat{\psi}$ unchanged. For example, within the Lognormal family with fixed shape $\sigma = 1.0$, increasing e^{μ} from 1 to 5 multiplies the underlying improvements by 5; correspondingly, the estimated cutoff increases from $\hat{c} \approx 0.856$ to $\hat{c} \approx 4.288$ (a $5\times$ change), while $\hat{\psi}$ remains 0.736 at both endpoints.

We estimate the threshold $\hat{c}$ and patience scalar $\hat{\psi}$. For each distribution parameterization, we generate $N = 10^6$ pairs of draws to estimate the threshold $\hat{c} = \frac{1}{N}\sum(X_i' - X_i)^+$. We then do an additional $N = 10^6$ independent simulations using this threshold to empirically measure $\hat{\psi}$.

Table 3: Verification of scale invariance (fixed shape, varying scale). Empirical estimates $(\hat{c}, \hat{\psi})$ are derived from 10^6 Monte Carlo simulations.

Distribution	**Shape (Fixed)**	**Scale Param**	**Threshold $\hat{c}$**	**Patience $\hat{\psi}$**
Exponential	$k = 1.0$	$b = 0.1 \to 10.0$	$0.05 \to 5.002$	**0.697**
Half-Normal	N/A	$\sigma = 0.2 \to 20.0$	$0.066 \to 6.62$	**0.668**
Gamma	$k = 2.0$	$\theta = 0.5 \to 8.0$	$0.374 \to 6.001$	**0.675**
Lognormal	$\sigma = 1.0$	$e^{\mu} = 1 \to 5$	$0.856 \to 4.288$	**0.736**

Shape dependence. For some families, changing the "shape" parameter also changes the overall level of X (e.g., Gamma with fixed θ changes mean $k\theta$), which would move the cutoff $c(\mathcal{D}) = \mathbb{E}[(X' - X)^+]$ for unit reasons rather than tail behavior. To ensure Table 4 reflects *shape* rather than scale, we normalize each family to have $\mathbb{E}[X] = 1$ for every shape value. Specifically, for Gamma we set $\theta = 1/k$; for Weibull we set the scale $\lambda = 1/\Gamma(1 + 1/c)$ (numerically $\lambda \approx 0.883$ for $c = 0.8$ and $\lambda \approx 1.108$ for $c = 1.5$); and for Lognormal we set $\mu = -\sigma^2/2$ (so $\mu = -1.125$ for $\sigma = 1.5$ and $\mu = -0.125$ for $\sigma = 0.5$).

With mean fixed, heavier tails increase both the cutoff and the agent's patience. For the Weibull family, moving from a lighter tail ($c = 1.5$) to a heavier tail ($c = 0.8$) raises the cutoff from $\hat{c} \approx 0.371$ to $\hat{c} \approx 0.581$, and raises the unconditional patience scalar from $\hat{\psi} \approx 0.668$ to $\hat{\psi} \approx 0.721$. This directly changes how often discretion is exercised early: the probability of spending in period 1 drops from $1 - \hat{\psi} \approx 0.332$ to $1 - \hat{\psi} \approx 0.279$ (about a 16% reduction in early overrides). The behavioral implication is visible at the decision level: a realized gain of $\Delta_1 = 0.5$ triggers an override under the lighter-tailed case ($0.5 > 0.371$) but is deferred under the heavier-tailed case ($0.5 < 0.581$), reflecting the option value of waiting for rarer opportunities.

Extreme regimes. Table 5 reports selected extreme configurations. The Lognormal and Pareto examples have high estimated patience. The near-degenerate Gamma example gives $\hat{\psi} \approx 0.655$. Note that this value is specific to that example, not a universal floor for near-degenerate distributions.

Table 4: Verification of shape dependence (varying shape). Empirical estimates $(\hat{c}, \hat{\psi})$ are derived from 10^6 Monte Carlo simulations.

Distribution	Shape Param	Shape Description	Threshold $\hat{c}$	Patience $\hat{\psi}$
Gamma	$k = 0.5$	L-shaped	0.637	**0.732**
Gamma	$k = 2.0$	Unimodal	0.375	0.675
Weibull	$c = 0.8$	Stretched-exponential tail	0.581	**0.721**
Weibull	$c = 1.5$	Faster-than-exponential tail	0.371	0.668
Lognormal	$\sigma = 1.5$	Larger log-scale spread	0.711	**0.802**
Lognormal	$\sigma = 0.5$	Smaller log-scale spread	0.276	0.678

Table 5: Extreme regimes: empirical $\hat{\psi}$ from 10^6 Monte Carlo draws. Excess kurtosis is a finite-sample statistic used as a proxy for tail heaviness (Normal: $\gamma_2 = 0$; larger values indicate heavier tails when the fourth moment is finite [28]).

Distribution	Tail Behavior	Excess Kurtosis	Patience $\hat{\psi}$
Lognormal ($\sigma = 5$)	Extreme Heavy	$\approx 8 \times 10^5$	**0.994**
Pareto ($b = 1.05$)	Power Law	$\approx 6 \times 10^5$	**0.962**
Gamma ($k = 10^4$)	Near-degenerate Gamma	≈ 0	**0.655**

The extreme-tail Monte Carlo estimates can be unstable, and should be interpreted as such: Pareto with shape 1.05 has infinite variance, and the Lognormal with $\sigma = 5$ is highly variable.

E.8 Behavioral implications of shape

The patience scalar ψ determines the agent's observable behavior relative to two benchmarks: a full-information oracle and the agent's own across-period allocation. The following lemmas make this precise.

Lemma E.1 (DP vs. Oracle Benchmark, two-period)**.** *Let π^* be the optimal DP policy and $\tilde{\pi}$ the oracle policy with perfect foresight of (Δ_1, Δ_2).*

1. *At $t = 1$, the DP policy is more conservative (less likely to spend) than the oracle if and only if $F_{\Delta|\mathsf{A}_t^c}(\mathcal{T}_{2,1}) > \frac{1-p}{2}$.*
2. *At $t = 2$, the DP policy is more aggressive (more likely to spend) than the oracle if and only if $F_{\Delta|\mathsf{A}_t^c}(\mathcal{T}_{2,1}) > \frac{1}{2}$.*

Proof of Lemma E.1. The proof proceeds by direct comparison of the probabilities derived in Equations (20), (21), and (22).

Part 1 (Comparison at $t = 1$): The DP policy is more conservative than the oracle at $t = 1$ when:

$$\begin{aligned} \overset{\pi^*}{\Pr}(\text{spend at } 1 \mid \mathsf{A}_1^c) &< \overset{\tilde{\pi}}{\Pr}(\text{spend at } 1 \mid \mathsf{A}_1^c) \\ 1 - F_{\Delta|\mathsf{A}_t^c}(\mathcal{T}_{2,1}) &< \frac{1+p}{2} \\ \frac{1-p}{2} &< F_{\Delta|\mathsf{A}_t^c}(\mathcal{T}_{2,1}). \end{aligned}$$

The condition is reversed if the DP policy is more aggressive.

Part 2 (Comparison at $t = 2$): The DP policy is more aggressive than the oracle at $t = 2$ when:

$$\begin{aligned}
\overset{\pi^*}{\Pr}(\text{spend at } 2 \mid \mathsf{A}_2^c) &> \overset{\tilde{\pi}}{\Pr}(\text{spend at } 2 \mid \mathsf{A}_2^c) \\
p + (1-p)F_{\Delta|\mathsf{A}_t^c}(\mathcal{T}_{2,1}) &> \frac{1+p}{2} \\
(1-p)F_{\Delta|\mathsf{A}_t^c}(\mathcal{T}_{2,1}) &> \frac{1+p}{2} - p \\
(1-p)F_{\Delta|\mathsf{A}_t^c}(\mathcal{T}_{2,1}) &> \frac{1-p}{2} \\
F_{\Delta|\mathsf{A}_t^c}(\mathcal{T}_{2,1}) &> \frac{1}{2} \quad (\text{since } 1-p > 0).
\end{aligned}$$

The condition is reversed if the DP policy is more conservative. □

Lemma E.2 (Internal DP Behavior, two-period)**.** *The optimal DP policy π^* is more likely to spend its budget at $t = 2$ than at $t = 1$ (conditional on misalignment in each period) if and only if $F_{\Delta|\mathsf{A}_t^c}(\mathcal{T}_{2,1}) > \frac{1-p}{2-p}$.*

Proof of Lemma E.2. We compare the conditional probability of spending at $t = 2$ with that at $t = 1$. The DP is more likely to spend at $t = 2$ if:

$$\begin{aligned}
\overset{\pi^*}{\Pr}(\text{spend at } 2 \mid \mathsf{A}_2^c) &> \overset{\pi^*}{\Pr}(\text{spend at } 1 \mid \mathsf{A}_1^c) \\
p + (1-p)F_{\Delta|\mathsf{A}_t^c}(\mathcal{T}_{2,1}) &> 1 - F_{\Delta|\mathsf{A}_t^c}(\mathcal{T}_{2,1}) \\
F_{\Delta|\mathsf{A}_t^c}(\mathcal{T}_{2,1}) + (1-p)F_{\Delta|\mathsf{A}_t^c}(\mathcal{T}_{2,1}) &> 1 - p \\
(1+1-p)F_{\Delta|\mathsf{A}_t^c}(\mathcal{T}_{2,1}) &> 1 - p \\
(2-p)F_{\Delta|\mathsf{A}_t^c}(\mathcal{T}_{2,1}) &> 1 - p \\
F_{\Delta|\mathsf{A}_t^c}(\mathcal{T}_{2,1}) &> \frac{1-p}{2-p} \quad (\text{since } 2-p > 0).
\end{aligned}$$

This shows that the DP defers spending to the second period when the cumulative mass of "small" gains below the threshold $\mathcal{T}_{2,1}$ is sufficiently large. The inequality is reversed if the DP is more likely to spend early. □

Since $\psi = p + (1-p)F_{\Delta|\mathsf{A}_t^c}(\mathcal{T}_{2,1})$ (Appendix E.3), the condition in Lemma E.2 is equivalent to $\psi > \frac{1}{2-p}$, and the condition in part 1 of Lemma E.1 is equivalent to $\psi > \frac{1+p^2}{2}$. The two lemmas combine into a single comparative-statics statement linking tail thickness to patience:

Corollary E.1 (Tail thickness and patience)**.** *Consider a family of base shapes $\mathcal{D}_\eta$ where the parameter η increases tail thickness. If $\psi(\mathcal{D}_\eta)$ is increasing in η, then for a fixed alignment probability p:*

- *As ψ exceeds $\frac{1}{2-p}$, the DP becomes internally patient, spending more often at $t = 2$ than at $t = 1$ (Lemma E.2).*
- *As ψ exceeds $\frac{(1+p^2)}{2}$, the DP becomes more conservative than the full-information oracle at $t = 1$ (Lemma E.1).*

F Computing per-period spending profiles

This appendix details the computation of the observable spending curve for the general (T, K) problem, as summarized in Section 3.4.

F.1 Conditional spending probability

For a fixed state (τ, k), the agent spends if a misalignment occurs *and* the gain clears the threshold:

$$q_{\tau,k} = \Pr(\text{spend} \mid \tau, k) = (1-p) \cdot \Pr(\Delta > \mathcal{T}_{\tau,k} \mid \mathsf{A}_t^c).$$

We show this equals $q_{\tau,k} = 1 - F_\Delta(\mathcal{T}_{\tau,k})$. The unconditional CDF of Δ at $x \geq 0$ is:

$$F_\Delta(x) = \Pr(\Delta \leq x) = p \cdot \underbrace{\Pr(\Delta \leq x \mid \mathsf{A}_t)}_{=1 \text{ since } \Delta=0 \text{ under } \mathsf{A}_t} + (1-p) \cdot \Pr(\Delta \leq x \mid \mathsf{A}_t^c) = p + (1-p) \cdot F_{\Delta|\mathsf{A}_t^c}(x).$$

Therefore $1 - F_\Delta(x) = (1-p)(1 - F_{\Delta|\mathsf{A}_t^c}(x)) = (1-p) \cdot \Pr(\Delta > x \mid \mathsf{A}_t^c)$, which gives $q_{\tau,k} = 1 - F_\Delta(\mathcal{T}_{\tau,k})$. By Theorem 3.1(iii), this probability is invariant to scale and location: $q_{\tau,k} = 1 - G_{\mathcal{D}}(\theta_{\tau,k}(\mathcal{D}))$.

F.2 Evolution of budget distribution as described in Section 3.4

Since spending is stochastic, the remaining budget k at time t is a random variable. Let $\pi_t(k)$ be the probability that the agent has exactly k budget units remaining at the start of period t. We compute this via a forward pass:

Initialize. $\pi_1(K) = 1$, and $\pi_1(k) = 0$ for $k < K$.

Forward update. For $t = 1, \ldots, T-1$, let $\tau = T - t + 1$. The budget evolves as:

$$\pi_{t+1}(k) = \underbrace{\pi_t(k)\big(1 - q_{\tau,k}\big)}_{\text{had } k\text{, didn't spend}} + \underbrace{\pi_t(k+1) q_{\tau,k+1}}_{\text{had } k+1\text{, spent}}. \tag{25}$$

Boundary terms. $\pi_t(K+1) = 0$, $q_{\tau,0} = 0$ (no budget), and $q_{\tau,k \geq \tau} = 1 - p$ (ample budget, $\mathcal{T}_{\tau,k} = 0$).

Aggregate spending profile. The unconditional probability of overriding at time t is the expectation over budget states:

$$\Pr(\text{spend at time } t) = \sum_{k=1}^{K} \pi_t(k) \cdot q_{\tau,k}. \tag{26}$$

This yields the spending profile $\big(\Pr(\text{spend at } 1), \ldots, \Pr(\text{spend at } T)\big)$, the observable "spending curve" of the optimal policy. Since each $q_{\tau,k}$ is scale- and location-free (Theorem 3.1), and the budget distribution $\pi_t(k)$ is computed from these probabilities via the deterministic update (25), the entire spending profile inherits the same invariance. This T-vector is the natural generalization of the two-period scalar ψ: it captures all observable discretionary behavior and depends only on $\mathcal{D}$, p, T, and K.

F.3 Threshold visualization

Figure 2 shows the optimal threshold $\mathcal{T}_{\tau,k}$ at every state (τ, k) for Pareto-distributed gains with $T = 20$ and $K = 5$.

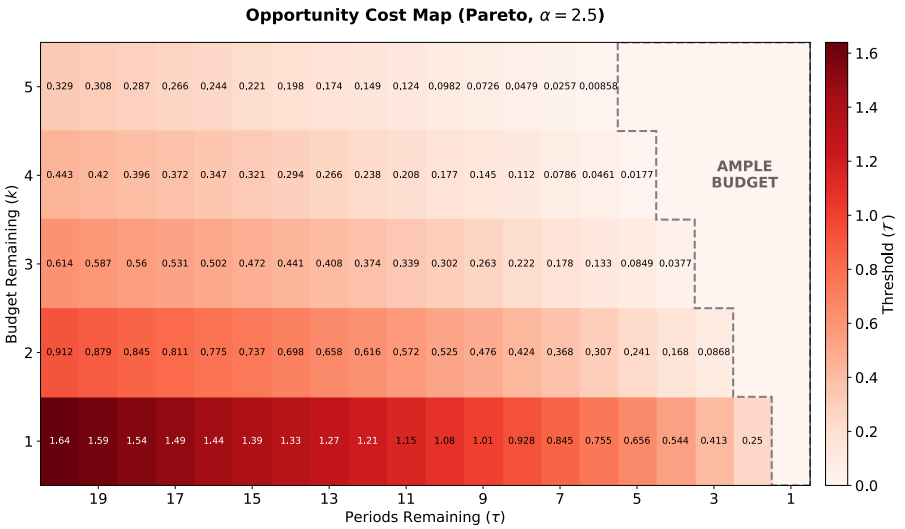


Figure 2: The opportunity cost map: heatmap of optimal thresholds $\mathcal{T}_{\tau,k}$ for $(T = 20, K = 5)$ with a Pareto distribution scaled to mean 1. Darker regions indicate states where the agent requires a very high gain to justify spending. The thresholds naturally decay as time runs out ($\tau \to 0$) or budget becomes abundant relative to time ($k \to \tau$).

G Empirical analysis: full details

To complement our theoretical characterization of threshold-based override behavior, we examine placement records from a Homeless Management Information System (HMIS) to test whether discretionary deviations from a recovered baseline policy vary systematically with operational constraints. We ask two questions. First, how does discretionary behavior respond to shifts in the operational constraints that bind the caseworker? And second, how is the use of discretion spread over the calendar?

Computing infrastructure used: All experiments and analysis were conducted on a 16-inch MacBook Pro (2024). The machine was equipped with an Apple M4 Pro system on a chip (SoC), featuring a 14-core CPU, a 20-core GPU, and a 16-core Neural Engine. It included 24 GB of unified memory and a 512 GB solid-state drive (SSD), running macOS 15.

G.1 Operational setting: agent, policy, and environment

We map our theoretical framework to the context of homeless service allocation. The caseworker is the decision-maker, and the allocation setting presents the same temporal tradeoff formalized in Section 2: Transitional Housing (TH) serves as the scarce, resource-intensive intervention, while Emergency Shelter (ES) serves as the default, lower-intensity baseline. Caseworkers possess the discretionary authority to override the default assignment to place a client in scarce transitional housing or, alternatively, move clients from scarce housing to the more readily available option (ES in this case).

To measure this discretion, we must first define the "Policy" from which the agent deviates. Drawing on recent work by Pokharel et al. [33], we operationalize the default policy not as a rigid bureaucratic manual—which did not exist in this historical period—but as the dominant set of heuristic rules governing the system. Street-level bureaucracy theory posits that agents operate under cognitive constraints by developing simplified heuristics to process routine cases quickly [25]. We model these heuristics using the Short Explainable Rules - Decision Tree (SER-DT) algorithm [39]. SER-DT explicitly optimizes for interpretability by penalizing complexity, producing short, depth-constrained decision trees (we set the maximum depth to 4) that capture systematic variance in assignments. We treat the prediction P_i of these trees as the observable baseline: it represents the decision a "typical" bureaucrat in this time-period would make. Thus, using the same definition of discretion as Pokharel

et al. [33], a discretionary event is defined as any instance in which the caseworker's actual decision differs from this baseline ($A_i \neq P_i$).

G.2 Data description and filtering

We analyze administrative records from the St. Louis Homeless Management Information System (HMIS), covering households assessed between 2008 and 2014. The data includes records of service entry and exit, linked to detailed household characteristics.

Time context (the role of discretion). The study period (2008–2014) represents a unique window for observing natural discretionary behavior. During this time, the homeless service system lacked more explicit, formulaic prioritization mandates (such as the VI-SPDAT [29] adopted later). Instead, caseworkers relied on professional judgment to balance household vulnerability against expected outcomes.

Feature selection. To train the SER-DT policy model described above, we utilize 34 household features available upon request for services, such as age, gender, income, disability status, and family composition. These features represent the information set available to the agent when forming the baseline heuristic rules (P_i) (the list of all features used is detailed in Table 6).

G.3 Operationalization of discretion

For each household i in the dataset, we observe two key variables: (i) the predicted assignment generated by the heuristic policy (P_i), and (ii) the caseworker's actual final placement decision (A_i).

We restrict our attention to the two primary intervention types: ES and TH. We exclude observations from other project types (i.e., rapid rehousing, prevention) to maintain consistency with the finite-budget resource allocation problem defined in Section 2, as ES and TH provide a clear hierarchy on the intensity of intervention assignment with TH being the more resource-intensive intervention. The restriction applies to the tree's recommendation as well as to the placement: overrides are counted only among households whose recommendation and placement are both ES or TH, and these households are the trials of the overall model. We then define a discretionary event as an instance in which the caseworker overrides the default policy recommendation ($A_i \neq P_i$). We classify these overrides into three indicators to capture the direction of resource usage:

- **Overall Discretion $\left(D_i^{\text{all}}\right)$:** An indicator that takes value 1 if the caseworker deviates from the policy in any direction.
$$D_i^{\text{all}} = \mathbb{I}\{P_i \neq A_i\}.$$
- **Upward Discretion / "Upgrading" $\left(D_i^{\uparrow}\right)$:** An indicator for instances where the policy recommends the default option (ES), but the caseworker "upgrades" the client to the scarce resource (TH).
$$D_i^{\uparrow} = \mathbb{I}\{P_i = \text{ES} \wedge A_i = \text{TH}\}.$$
In the language of our dynamic model, these are instances where the caseworker "spends" the discretionary budget to capture a perceived utility gain.
- **Conservative Discretion / "Rationing" $\left(D_i^{\downarrow}\right)$:** An indicator for instances where the policy recommends the scarce resource (TH), but the caseworker picks the 'default' option (ES).
$$D_i^{\downarrow} = \mathbb{I}\{P_i = \text{TH} \wedge A_i = \text{ES}\}.$$
These cases represent a refusal to allocate the more intensive resource even when the policy recommends it, effectively "saving" the budget for future periods.

Errors in the recovered baseline No record of the routine exists for this period, so some cases counted as overrides may be tree errors, and some real overrides may be missed. The tree sees only the 34 household features in Table 6, none of which are the calendar or capacity regressors of Appendix G.4. If its error rates do not move with those regressors, misclassification adds overrides on quiet days and hides them on busy ones, pulling the odds ratios toward one: tree error can weaken a real association but cannot create one. Error rates could move with the regressors only through the

case mix, for example if households arriving after a weekend differ in ways that change how often the tree makes a mistake. We assume the case mix is stable in this sense. We cannot test this, because no part of the 2008–2014 data records the routine independently, and for the same reason the share of measured overrides that are tree errors cannot be estimated. Two results suggest the recovered routine is reliable. Ten refits of the tree in Pokharel et al. [33] give nearly identical predictions (mean pairwise Spearman correlation 0.985 for TH and 0.890 for ES), and in our data XGBoost at its default settings predicts placements far better than the depth-4 trees for both ES and TH (Appendix G.7), so deviations from the short rule follow consistent patterns rather than noise.

From household decisions to daily binomial outcomes Our theoretical model characterizes optimal overrides through a dynamic threshold rule. Behaviorally, this appears as a time-varying probability of exercising discretion. Although the theory is formulated in abstract decision periods, the operational reality of the homeless services system–staffing schedules, intake workflows, and resource tracking–is organized by *calendar days*. To bridge this gap, we aggregate individual placement decisions to the daily level and treat each day as a binomial experiment, measuring the frequency of discretionary overrides under varying system conditions. Let t index calendar days and let $\mathcal{I}_t$ denote the set of households whose entry date falls on day t and whose realized placement lies in $\{\text{ES}, \text{TH}\}$. Define the daily volume $n_t = |\mathcal{I}_t|$, and the daily counts of discretionary overrides as:

$$Y_t^{\text{all}} = \sum_{i \in \mathcal{I}_t} D_i^{\text{all}}, \qquad Y_t^{\uparrow} = \sum_{i \in \mathcal{I}_t} D_i^{\uparrow}, \qquad Y_t^{\downarrow} = \sum_{i \in \mathcal{I}_t} D_i^{\downarrow}.$$

Reference sets for directional discretion. For directional discretion, the relevant denominator is not total daily volume, but the set of cases for which that direction is *feasible*. We refer to these as eligible *trial sets*. Concretely, Upgrading (ES→TH) is only possible for households whose baseline recommendation is ES, while Rationing (TH→ES) is only possible for households recommended for TH. Accordingly, we define daily trial-set sizes:

$$n_t^{\text{ES}} = \sum_{i \in \mathcal{I}_t} \mathbb{I}\{P_i = \text{ES}\}, \qquad n_t^{\text{TH}} = \sum_{i \in \mathcal{I}_t} \mathbb{I}\{P_i = \text{TH}\},$$

and we model daily override counts as binomial draws:

$$Y_t^{\text{all}} \sim \text{Binom}\left(n_t, p_t^{\text{all}}\right), \qquad Y_t^{\uparrow} \sim \text{Binom}\left(n_t^{\text{ES}}, p_t^{\uparrow}\right), \qquad Y_t^{\downarrow} \sim \text{Binom}\left(n_t^{\text{TH}}, p_t^{\downarrow}\right),$$

restricting to days with active caseloads ($n_t > 0$) and, for directional models, non-empty trial sets $\left(n_t^{\text{ES}} > 0 \text{ or } n_t^{\text{TH}} > 0\right)$.

G.4 Constructing operational state variables

To test the predictions of the Budgeted Discretion framework, we reconstruct daily proxies for the operational "state" visible to the agent. We engineer features intended to capture (i) workload and throughput, (ii) short-run slack in ES and TH, and (iii) temporal cycles. To preserve temporal ordering (and to avoid mechanical dependence on same-day decisions), all operational features are constructed using only information realized *strictly prior* to day t.

G.4.1 Rolling state (immediate operational conditions)

Our primary specification summarizes recent system conditions over a rolling 7-day window. For each day t, we compute total inflow (assignments) and outflow (exits) over $[t-7, t-1]$:

$$L_t^{\text{roll}} = \left[\text{Assign}_{t-7:t-1}^{\text{ES}}, \qquad \text{Assign}_{t-7:t-1}^{\text{TH}}, \qquad \text{Exit}_{t-7:t-1}^{\text{ES}}, \qquad \text{Exit}_{t-7:t-1}^{\text{TH}}\right]^{\top}.$$

Assignments proxy recent throughput and workflow burden. Exits proxy *openings* (short-run slack) for each intervention. When exits are high, recent capacity has been freed; when exits are low, the intervention is tighter. In the language of our model, these exit-based measures provide a reduced-form proxy for tightness and, thus, for the effective threshold at which discretion is exercised.

G.4.2 Weekly block state (administrative cycle robustness)

As a robustness check, we replace rolling totals with totals from the prior full calendar week (lagged by one full week):

$$L_t^{\text{block}} = \Big[\text{Assign}^{\text{ES}}_{\text{prev wk}}, \quad \text{Assign}^{\text{TH}}_{\text{prev wk}}, \quad \text{Exit}^{\text{ES}}_{\text{prev wk}}, \quad \text{Exit}^{\text{TH}}_{\text{prev wk}}\Big]^\top.$$

This block specification captures coarser batching and reporting cycles. The results are generally robust to this alternative state definition; see Appendix H.

G.4.3 Lagged TH composition

We also control for the prevailing ES/TH mix using the lagged (by one day) share of TH placements:

$$\text{THShare}_{t-1} = \frac{\text{TH}_{t-1}}{n_{t-1}}, \qquad \text{TH}_{t-1} = \sum_{i \in \mathcal{I}_{t-1}} \mathbb{I}\{A_i = \text{TH}\}. \tag{27}$$

Using the lagged value avoids mechanical correlation with same-day placements while absorbing slow-moving shifts in the homeless service system's placement regime.

G.4.4 Temporal structure

We include calendar controls, capturing systematic cycles in operations:

- **Day-type indicators:** to separate a start-of-week effect from weekend operations, we use a three-level indicator $\text{DayType}(t) \in \{\text{Mon}, \text{Tue–Fri}, \text{Weekend}\}$ with Tue–Fri as the reference.
- **Month-of-year fixed effects:** indicators for $\text{month}(t) \in \{1, \ldots, 12\}$ with October as the reference, since October serves as the start of the federal fiscal year.
- **Federal holidays:** an indicator H_t for U.S. federal holidays.

G.5 Model specification

We estimate the drivers of daily discretion using a generalized linear model (GLM) with a binomial family and logit link. For each outcome $j \in \{\text{all}, \uparrow, \downarrow\}$ and state specification $L_t \in \big\{L_t^{\text{roll}}, L_t^{\text{block}}\big\}$, we estimate:

$$\log\left(\frac{p_t^j}{1-p_t^j}\right) = \alpha^j_{\text{DayType}(t)} + \gamma^j_{\text{month}(t)} + \beta_S^j \cdot \text{THShare}_{t-1} + \beta_H^j \cdot H_t + (\beta_L^j)^\top L_t.$$

Observations are weighted by their respective trial counts: n_t for Overall Discretion, n_t^{ES} for Upgrading, and n_t^{TH} for Rationing. This ensures that days with higher relevant volumes contribute proportionally more to the likelihood.

G.6 Features used in SER-DT

Table 6: A comprehensive description of all the features that we used in the creation of the decision trees is shown in this table. All categorical variables were one-hot encoded.

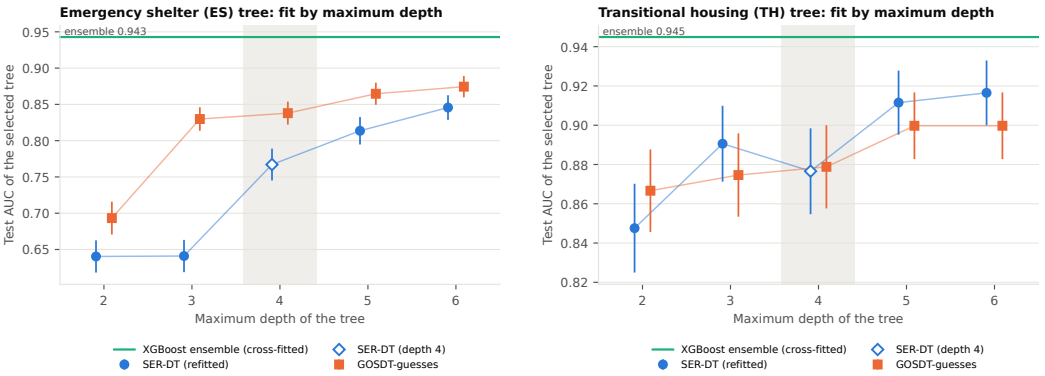
Binary features	
HUDChronicHomeless	Whether or not the client is chronically homeless.
Gender	The original documentation allows 8 values, but the data contain only Male and Female.
SpousePresent	Variable indicating whether or not client has a spouse.
Categorical features	
PrimaryRace	The primary race of client. Can be one of 7 categories.
Ethnicity	The ethnicity of client. Can be one of 5 categories.
PriorResidence	The type of residence client had before entering system. Can be one of 25 categories.
VeteranStatus	Variable indicating whether or not client is a veteran. Can take one of 5 values.
DisablingCondition	Variable indicating whether or not client has a disabling condition. Can be one of 3 categories.
ReceivePhysicalDisabilityServices	Variable indicating whether or not client received disability services. Can be one of 5 categories.
HasDevelopmentalDisability	Variable indicating whether or not client has developmental disability. Can be one of 5 categories.
ReceiveDevelopmentalDisabilityServices	Variable indicating whether or not client received services for developmental disability. Can be one of 5 categories.
HasChronicHealthCondition	Variable indicating whether or not client has a chronic health condition. Can be one of 5 categories.
ReceiveChronicHealthServices	Variable indicating whether or not client received services for a chronic health condition. Can be one of 5 categories.
HasHIVAIDS	Variable indicating whether or not client has HIV/AIDS. Can be one of 5 categories.
ReceiveHIVAIDSServices	Variable indicating whether or not client received services for HIV/AIDS. Can be one of 5 categories.
HasMentalHealthProblem	Variable indicating whether or not client has a mental health condition. Can be one of 5 categories.
ReceiveMentalHealthServices	Variable indicating whether or not client received services related to mental health. Can be one of 5 categories.
HasSubstanceAbuseProblem	Variable indicating whether or not client has a substance abuse problem. Can be one of 5 categories.
ReceiveSubstanceAbuseServices	Variable indicating whether or not client received services related to substance abuse. Can be one of 5 categories.
DomesticViolenceSurvivor	Variable indicating whether or not client is a survivor of domestic violence. Can be one of 5 categories.
Continuous features	
Age	The age of the client.
Calls	The number of calls made to hotline before entry.
Wait	The time (in days) elapsed since first call to hotline to intervention assignment.
RatioOfNumCallstoWaitTime	Ratio of number of calls to wait time. Calculated variable to measure 'persistence'.
MonthlyAmount	The total income (in dollars) of client per month.
numMembers	The number of members in the client's household.
Children	The number of children in client's household.
Children 0-2	The number of children aged 0 to 2 in client's household.
Children 3-5	The number of children aged 3 to 5 in client's household.
Children 6-10	The number of children aged 6 to 10 in client's household.
Children 11-14	The number of children aged 11 to 14 in client's household.
Children 15-17	The number of children aged 15 to 17 in client's household.
UnrelatedChildren	The number of children in client's household not related to them.
UnrelatedAdults	The number of adults in client's household not related to them.

G.7 Sensitivity to tree depths

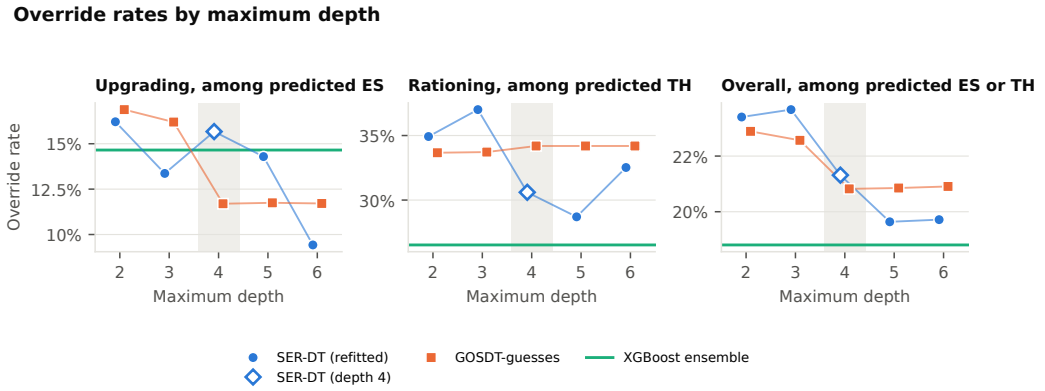
The complexity of the tree determines which departures from it are counted as discretion. Street-level bureaucracy theory describes routine behavior in terms of a small set of simple rules [25]. A more flexible model can incorporate systematic departures from those rules into the recovered routine, leaving fewer placements classified as overrides. The maximum tree depth therefore forms part of our definition of discretion. Here, we ask whether the empirical conclusions of Section 4 hold across alternative depth limits and fitting methods. We repeat the analysis with the short explainable decision trees (SER-DT) used in the main analysis [39], fitted at maximum depths 2, 3, 5, and 6. For predictive fit, override rates, and random-group tests, we also fit generalized and scalable optimal sparse decision trees [24] using the guessing strategies of McTavish et al. [27] (GOSDT-guesses) at depths 2 through 6. An XGBoost ensemble [7] at its default settings, one model per intervention with 100 trees of depth at most 6 and a learning rate of 0.3, provides a more flexible benchmark for these comparisons.

Our reference baseline consists of the depth-4 SER-DT trees from Pokharel et al. [33], with one tree for each intervention: emergency shelter (ES), transitional housing (TH), rapid rehousing (RRH), and homelessness prevention. The tables in the main body use predictions from these trees. At depth 4, we report their area under the receiver operating characteristic curve (AUC) and random-group percentiles, and use their predictions in the daily models reported in Section 4.

In keeping with the protocol of Pokharel et al. [33], we randomly split the households ten times, allocating 70% to fitting and 15% each to test and validation sets. For each intervention and depth, we retain the tree with the highest test accuracy across the ten splits. GOSDT-guesses uses the same splits, features, and selection criterion, with a penalty of 0.001 per leaf and candidate thresholds guessed by 40 boosted stumps. XGBoost uses five-fold cross-fitting. For the daily analyses, the overall model uses the specification in Table 1, which counts a day's overrides among placements predicted ES or TH. The depth-4 upgrading and rationing models are those in Tables 8 and 9.



(a) (b)



(c)

Figure 3: Predictive fit and measured discretion by maximum tree depth. Vertical axes are not shared: panels (a) and (b) use different scales, and each of the three panels in (c) has its own scale. (a, b) Test AUCs for the selected ES and TH trees, with DeLong 95% confidence intervals. Horizontal lines show XGBoost's cross-fitted AUCs. (c) Upgrading rates among placements predicted ES, rationing rates among placements predicted TH, and overall override rates among placements predicted ES or TH. Horizontal lines show the corresponding rates for XGBoost. Open diamonds denote the depth-4 SER-DT used in the main analysis. Filled circles denote SER-DT refits at the other depths, and squares denote GOSDT-guesses.

Predictive fit. Figure 3a–b reports test AUCs for the selected ES and TH trees, with 95% confidence intervals computed using DeLong's method. The ensemble's cross-fitted AUC provides a benchmark. The AUCs quoted in this paragraph and the next are means of the test AUC over the ten splits; the figure plots the test AUC of the selected tree, which is higher and varies more across depths. Predictive performance generally improves with depth, although the gains taper at different points across interventions and fitting methods. For ES, mean test AUC across the ten SER-DT splits rises from 0.638 at depth 2 to 0.827 at depth 6, with an increase of 0.030 between depths 5 and 6. Under GOSDT-guesses, ES AUC rises from 0.679 to 0.858, but gains fall below 0.01 per additional level from depth 5 onward. SER-DT depth 4, the tree used in the main analysis, has an AUC of 0.761, compared with 0.943 for XGBoost.

For TH, the gains taper earlier. SER-DT AUC rises from 0.837 at depth 2 to 0.908 at depth 6, with gains below 0.01 beyond depth 5; the depth-4 tree used in the main analysis has an AUC of 0.870. GOSDT-guesses AUC ranges from 0.855 to 0.873, with no increase of 0.01 or more between successive depths, while XGBoost achieves an AUC of 0.945. The improvements in fit must be weighed against the complexity of the resulting rules. The SER-DT trees reach their depth limits at every depth, growing from 3 leaves at depth 2 to as many as 38 at depth 6. GOSDT-guesses produces trees with 4 to 10 leaves, smaller than the SER-DT trees at depths 5 and 6, and reaches the

depth limit in nine of the ten ES and TH fits. We use depth 4 as our main specification to strike a reasonable balance between predictive performance and interpretability. Deeper trees can improve fit, particularly for ES, but our aim is to recover a compact set of rules that can represent routine decision-making. The analyses below examine whether the substantive conclusions depend on this choice.

Measured discretion. A more flexible baseline should account for more of the systematic variation in placements, reducing the share classified as overrides. Figure 3c broadly supports this expectation. Among placements predicted ES or TH, the overall override rate falls from 23.4% at depth 2 to 19.7% at depth 6 under SER-DT, compared with 21.3% for the original depth-4 baseline. Under GOSDT-guesses, the rate falls from 22.9% to 20.9%. XGBoost yields the lowest rate, at 18.8%. The decline is not monotonic: the SER-DT override rate increases between depths 2 and 3.

Most of the decline comes from upgrading. Its rate falls from 16.2% to 9.4% under SER-DT and from 16.9% to 11.7% under GOSDT-guesses. Rationing shows no comparable trend. Its rate ranges from 28.7% to 37.0% under SER-DT and remains between 33.7% and 34.2% across all GOSDT-guesses depths. XGBoost has the lowest rationing rate, at 26.5%, but its upgrading rate of 14.7% exceeds that of the deeper GOSDT-guesses trees. Baseline flexibility thus changes the amount of measured discretion, particularly upgrading.

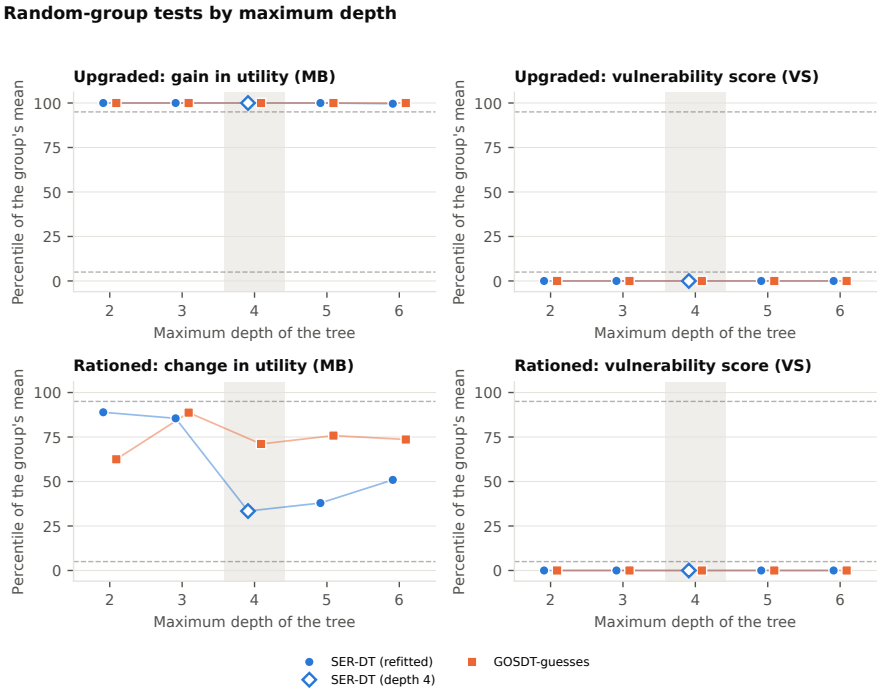


Figure 4: Random-group comparisons across tree depths. Each panel reports the percentile of an override group's mean among the means of 1,000 random groups with the same size and predicted intervention. Rows distinguish upgraded households (top) from rationed households (bottom); columns report utility gain (MB, left) and vulnerability score (VS, right). Open diamonds denote the depth-4 SER-DT baseline.

Random-group tests. We next examine whether changing the models alters the characteristics of households classified as receiving an override. Following Pokharel et al. [33], we compare upgraded and rationed households on two measures: the utility gain from the switch they received (MB in the figures) and the vulnerability score (VS). For each group, we draw 1,000 random groups of the same size from households with the same predicted intervention. We then locate the observed group's mean within the distribution of random-group means. The reported percentile is the share of random means below the observed mean. Upgraded households are those predicted ES and placed in TH; rationed households are those predicted TH and placed in ES.

The original analysis found that upgraded households had higher utility gains and lower vulnerability scores than the random groups. Rationed households also had lower vulnerability scores, but their utility gains did not differ significantly. All tested specifications reproduce this pattern (Figure 4). For the alternatives—four SER-DT depths, five GOSDT-guesses depths—upgraded households lie between the 99.6th and 100th percentiles for utility gain and at percentile 0.0 for vulnerability. Rationed households also lie at percentile 0.0 for vulnerability. Their utility-gain percentiles range from 37.9 to 88.9, with no statistically significant difference under any specification.

Daily models. Finally, we re-estimate the three daily models in Table 1 using overrides defined by SER-DT at depths 2, 3, 5, and 6. We first consider the ten odds ratios discussed in Section 4: ES and TH exits in the upgrading and rationing models; Monday, the weekend, November, December, and June in the overall model; and the weekend in the rationing model. At every alternative depth, each estimate lies within the corresponding 95% confidence interval for the depth-4 tree results reported in Table 1, covering all 40 comparisons. This also holds for the full set of 27 odds ratios: all 108 estimates lie within the depth-4 intervals. With the weekly block features in Tables 7–9, 107 of 108 estimates fall within the corresponding depth-4 intervals. The sole exception is the coefficient on ES exits in the upgrading model at depth 6. Its odds ratio is 0.945, just 0.002 below the depth-4 interval's lower bound of 0.947, placing the estimate slightly farther from one.

H Generalized linear model results

This appendix reports every coefficient of the daily models of Appendix G.5, for both the weekly block features and the rolling features of Appendix G.4. Each table gives the odds ratio of each term, its 95% confidence interval and its p-value; p-values below 0.05 are in bold, except for the intercept.

H.1 Weekly cycles

Tables 7, 8 and 9 report the models of Table 1 in full, with the day-type indicator (Monday, Tuesday to Friday, weekend), whose reference is Tuesday to Friday.

Table 7: Results of fitting the GLM for overall discretion, isolating weekly effects. The outcome counts each day's overrides among its placements predicted ES or TH.

	Block features			Rolling features		
Variable	Odds ratio	95% CI	p-value	Odds ratio	95% CI	p-value
Intercept and resource availability						
Intercept	0.41708	[0.30159, 0.57681]	1.25E-07	0.36135	[0.26132, 0.49966]	7.46E-10
THShare_{t-1}	1.0854	[0.88976, 1.3241]	0.41895	1.06	[0.859, 1.308]	0.58698
Holiday	0.86775	[0.54688, 1.3769]	0.54706	0.85699	[0.54011, 1.3598]	0.51235
ES assignments	0.99535	[0.98258, 1.0083]	0.4793	0.99656	[0.98381, 1.0095]	0.59987
TH assignments	0.98944	[0.97354, 1.0056]	0.1992	1.0059	[0.98902, 1.023]	0.49728
ES exits	0.98813	[0.9751, 1.0013]	0.078121	0.99468	[0.98201, 1.0075]	0.41477
per SD	0.9362	[0.86999, 1.0074]	0.078121	0.97004	[0.90166, 1.0436]	0.41477
TH exits	1.0082	[0.9914, 1.0252]	0.34206	0.99497	[0.97837, 1.0118]	0.55656
per SD	1.0366	[0.96252, 1.1164]	0.34206	0.97803	[0.90824, 1.0532]	0.55656
Weekly cycles (reference: Tue–Fri)						
Mon	1.2412	[1.0726, 1.4364]	**0.0037185**	1.2437	[1.0745, 1.4395]	**0.0034714**
Weekend	0.75872	[0.60177, 0.9566]	**0.019531**	0.76204	[0.60454, 0.96058]	**0.021427**
Month (reference: October)						
Nov	0.71122	[0.51642, 0.9795]	**0.036903**	0.70467	[0.51214, 0.96958]	**0.031578**
Dec	0.59141	[0.42812, 0.81698]	**0.0014417**	0.59163	[0.42812, 0.8176]	**0.0014723**
Jan	0.79854	[0.59368, 1.0741]	0.13691	0.78886	[0.58531, 1.0632]	0.11935
Feb	0.949	[0.70558, 1.2764]	0.7292	0.92882	[0.69123, 1.2481]	0.62426
Mar	0.842	[0.62609, 1.1324]	0.25527	0.81923	[0.60926, 1.1015]	0.1869
Apr	0.87372	[0.6481, 1.1779]	0.37575	0.87013	[0.64531, 1.1733]	0.36165
May	0.91588	[0.68211, 1.2298]	0.55895	0.9133	[0.68053, 1.2257]	0.5457
Jun	0.74992	[0.55261, 1.0177]	0.064662	0.71005	[0.52386, 0.96241]	**0.027318**
Jul	0.90081	[0.67094, 1.2094]	0.48711	0.89886	[0.67022, 1.2055]	0.47648
Aug	0.81364	[0.60183, 1.1]	0.18008	0.80124	[0.59337, 1.0819]	0.14815
Sep	0.8778	[0.65028, 1.1849]	0.3945	0.88483	[0.65584, 1.1938]	0.42327

Table 8: Results of fitting the GLM for upgrading, isolating weekly effects.

	Block features			Rolling features		
Variable	Odds ratio	95% CI	p-value	Odds ratio	95% CI	p-value
Intercept and resource availability						
Intercept	0.2536	[0.15612, 0.41192]	2.97E-08	0.21156	[0.12995, 0.34442]	4.20E-10
THShare_{t-1}	1.0524	[0.7859, 1.4093]	0.73164	0.93123	[0.68255, 1.2705]	0.65306
Holiday	1.0109	[0.52646, 1.9413]	0.97393	1.0059	[0.52407, 1.9308]	0.98587
ES assignments	0.99977	[0.98182, 1.018]	0.97977	0.99762	[0.97962, 1.016]	0.7977
TH assignments	0.99885	[0.97627, 1.0219]	0.92144	1.0208	[0.99684, 1.0453]	0.089403
ES exits	0.96525	[0.94666, 0.9842]	**0.00036287**	0.97053	[0.95258, 0.98882]	**0.0016862**
per SD	0.82254	[0.7388, 0.91579]	**0.00036287**	0.84319	[0.75805, 0.93789]	**0.0016862**
TH exits	1.0202	[0.99656, 1.0445]	0.094423	1.0212	[0.99764, 1.0454]	0.078135
per SD	1.0927	[0.98488, 1.2123]	0.094423	1.0968	[0.98965, 1.2156]	0.078135
Weekly cycles (reference: Tue–Fri)						
Mon	1.1256	[0.90662, 1.3974]	0.28381	1.1402	[0.91861, 1.4152]	0.23406
Weekend	0.78578	[0.55433, 1.1139]	0.17567	0.81324	[0.57366, 1.1529]	0.24566
Month (reference: October)						
Nov	0.91139	[0.56273, 1.4761]	0.70607	0.88028	[0.54414, 1.4241]	0.60338
Dec	0.74647	[0.45838, 1.2156]	0.23992	0.72972	[0.44763, 1.1896]	0.20632
Jan	1.2291	[0.7894, 1.9138]	0.36115	1.125	[0.71977, 1.7583]	0.60524
Feb	1.274	[0.81498, 1.9915]	0.28808	1.1875	[0.76002, 1.8554]	0.45042
Mar	1.1528	[0.73337, 1.812]	0.53781	1.0625	[0.67585, 1.6703]	0.79289
Apr	1.1349	[0.72185, 1.7842]	0.58362	1.0611	[0.67405, 1.6703]	0.79787
May	1.1371	[0.72388, 1.7863]	0.57705	1.1099	[0.70692, 1.7426]	0.65053
Jun	1.1223	[0.70555, 1.7852]	0.62611	1.0266	[0.64676, 1.6296]	0.91127
Jul	1.3444	[0.86745, 2.0835]	0.18556	1.2843	[0.82956, 1.9883]	0.26185
Aug	0.97605	[0.61305, 1.554]	0.91864	0.95199	[0.59898, 1.513]	0.83511
Sep	0.96196	[0.60527, 1.5288]	0.86966	0.93435	[0.5879, 1.485]	0.77392

Table 9: Results of fitting the GLM for rationing, isolating weekly effects.

	Block features			Rolling features		
Variable	Odds ratio	95% CI	p-value	Odds ratio	95% CI	p-value
Intercept and resource availability						
Intercept	0.74218	[0.46469, 1.1854]	0.21198	0.65816	[0.41327, 1.0482]	0.078095
THShare_{t-1}	1.0518	[0.79233, 1.3961]	0.72698	1.1504	[0.85331, 1.551]	0.35795
Holiday	0.79096	[0.40205, 1.5561]	0.49698	0.76989	[0.38982, 1.5205]	0.45136
ES assignments	0.99698	[0.97782, 1.0165]	0.75995	1.004	[0.98483, 1.0235]	0.68639
TH assignments	0.97263	[0.94968, 0.99612]	**0.022679**	0.98264	[0.95812, 1.0078]	0.17449
ES exits	1.0156	[0.99651, 1.035]	0.10985	1.0259	[1.0069, 1.0452]	**0.0073899**
per SD	1.0891	[0.9809, 1.2093]	0.10985	1.1569	[1.0399, 1.287]	**0.0073899**
TH exits	0.98491	[0.96048, 1.01]	0.23544	0.95233	[0.92806, 0.97724]	**0.00020875**
per SD	0.93493	[0.83658, 1.0448]	0.23544	0.8066	[0.71998, 0.90364]	**0.00020875**
Weekly cycles (reference: Tue–Fri)						
Mon	1.2687	[1.0305, 1.5621]	**0.024923**	1.2414	[1.0064, 1.5314]	**0.043472**
Weekend	0.63895	[0.46477, 0.87839]	**0.0058081**	0.63915	[0.46466, 0.87916]	**0.0059286**
Month (reference: October)						
Nov	0.66864	[0.42145, 1.0608]	0.087393	0.67713	[0.4265, 1.075]	0.098303
Dec	0.52615	[0.33201, 0.8338]	**0.0062624**	0.54741	[0.34478, 0.86912]	**0.010629**
Jan	0.55675	[0.36286, 0.85422]	**0.0073317**	0.6124	[0.3975, 0.94348]	**0.026162**
Feb	0.81131	[0.52933, 1.2435]	0.33719	0.82134	[0.53633, 1.2578]	0.36539
Mar	0.66791	[0.43974, 1.0145]	0.058415	0.6953	[0.45728, 1.0572]	0.089186
Apr	0.77213	[0.50313, 1.1849]	0.23665	0.82588	[0.5372, 1.2697]	0.3833
May	0.82495	[0.54317, 1.2529]	0.36677	0.84278	[0.55434, 1.2813]	0.42358
Jun	0.52548	[0.34164, 0.80823]	**0.0033985**	0.52213	[0.33891, 0.80442]	**0.0032097**
Jul	0.69954	[0.45494, 1.0756]	0.10356	0.72898	[0.474, 1.1211]	0.15005
Aug	0.72658	[0.47571, 1.1098]	0.13939	0.73416	[0.4804, 1.122]	0.15325
Sep	0.90581	[0.59132, 1.3876]	0.64939	0.96224	[0.62726, 1.4761]	0.86004

H.2 Full day-of-week configuration

Tables 10, 11 and 12 replace the day-type indicator with an indicator for each day of the week, whose reference is Monday.

Table 10: Results of fitting the GLM for overall discretion, full day-of-week configuration. The outcome counts each day's overrides among its placements predicted ES or TH.

	Block features			Rolling features		
Variable	Odds ratio	95% CI	p-value	Odds ratio	95% CI	p-value
Intercept and resource availability						
Intercept	0.51593	[0.36515, 0.72898]	0.00017519	0.44808	[0.31726, 0.63286]	5.19E-06
THShare_{t-1}	1.0656	[0.87294, 1.3008]	0.5323	1.0386	[0.84097, 1.2826]	0.72534
Holiday	0.86017	[0.54213, 1.3648]	0.52248	0.85026	[0.53589, 1.349]	0.49096
ES assignments	0.99487	[0.98209, 1.0078]	0.43526	0.99685	[0.98407, 1.0098]	0.6322
TH assignments	0.99075	[0.97479, 1.007]	0.26236	1.0068	[0.98993, 1.024]	0.43105
ES exits	0.98908	[0.976, 1.0023]	0.10564	0.99484	[0.98215, 1.0077]	0.43003
per SD	0.94114	[0.87446, 1.0129]	0.10564	0.97094	[0.90238, 1.0447]	0.43003
TH exits	1.0078	[0.991, 1.0249]	0.36539	0.99512	[0.97848, 1.0121]	0.5701
per SD	1.0349	[0.96079, 1.1148]	0.36539	0.97872	[0.9087, 1.0541]	0.5701
Day of the week (reference: Monday)						
Tue	0.91052	[0.75661, 1.0957]	0.32108	0.91066	[0.75667, 1.096]	0.32206
Wed	0.75151	[0.61915, 0.91216]	**0.0038516**	0.75019	[0.61799, 0.91068]	**0.003662**
Thu	0.6826	[0.55802, 0.835]	**0.00020411**	0.67821	[0.5543, 0.82983]	**0.00016193**
Fri	0.87231	[0.71823, 1.0594]	0.16829	0.8712	[0.71719, 1.0583]	0.16478
Sat	0.70326	[0.53921, 0.91721]	**0.0093869**	0.70706	[0.5422, 0.92206]	**0.010498**
Sun	0.26723	[0.13345, 0.53511]	**0.00019535**	0.26419	[0.132, 0.52879]	**0.00017016**
Month (reference: October)						
Nov	0.70704	[0.51298, 0.97452]	**0.034208**	0.70113	[0.50917, 0.96546]	**0.029604**
Dec	0.59249	[0.42865, 0.81894]	**0.0015274**	0.59113	[0.42753, 0.81734]	**0.0014723**
Jan	0.79185	[0.58846, 1.0655]	0.12335	0.78025	[0.57861, 1.0522]	0.10381
Feb	0.94672	[0.70356, 1.2739]	0.71772	0.92531	[0.6883, 1.2439]	0.60716
Mar	0.84117	[0.62512, 1.1319]	0.25346	0.81751	[0.6076, 1.0999]	0.18325
Apr	0.87086	[0.6457, 1.1746]	0.36499	0.86636	[0.64222, 1.1687]	0.34765
May	0.91587	[0.68173, 1.2304]	0.5596	0.91095	[0.67843, 1.2232]	0.53508
Jun	0.74448	[0.54845, 1.0106]	0.058424	0.70598	[0.52072, 0.95714]	**0.024957**
Jul	0.891	[0.66316, 1.1971]	0.44374	0.88837	[0.66189, 1.1923]	0.43049
Aug	0.80545	[0.59543, 1.0895]	0.16042	0.79331	[0.58717, 1.0718]	0.13151
Sep	0.8722	[0.64561, 1.1783]	0.37298	0.8768	[0.64937, 1.1839]	0.39079

Table 11: Results of fitting the GLM for upgrading, full day-of-week configuration.

	Block features			Rolling features		
Variable	Odds ratio	95% CI	p-value	Odds ratio	95% CI	p-value
Intercept and resource availability						
Intercept	0.28191	[0.16787, 0.47341]	1.69E-06	0.23865	[0.14182, 0.40162]	6.84E-08
THShare_{t-1}	1.0522	[0.7856, 1.4094]	0.7327	0.92917	[0.68095, 1.2679]	0.64314
Holiday	1.0059	[0.52395, 1.9312]	0.98589	1.0019	[0.52211, 1.9225]	0.99553
ES assignments	0.99998	[0.98202, 1.0183]	0.9986	0.99736	[0.97935, 1.0157]	0.77631
TH assignments	0.99794	[0.97532, 1.0211]	0.85997	1.0207	[0.99667, 1.0453]	0.091831
ES exits	0.96527	[0.94666, 0.98424]	**0.0003715**	0.97106	[0.95312, 0.98933]	**0.0020198**
per SD	0.82263	[0.73877, 0.916]	**0.0003715**	0.8458	[0.76049, 0.94068]	**0.0020198**
TH exits	1.021	[0.99726, 1.0453]	0.083309	1.0213	[0.99764, 1.0455]	0.077999
per SD	1.0964	[0.98795, 1.2167]	0.083309	1.0971	[0.98967, 1.2162]	0.077999
Day of the week (reference: Monday)						
Tue	0.96777	[0.73978, 1.266]	0.8111	0.96421	[0.7373, 1.261]	0.79009
Wed	0.73086	[0.54716, 0.97622]	**0.033763**	0.72047	[0.53954, 0.96207]	**0.026283**
Thu	0.83432	[0.6269, 1.1104]	0.21419	0.82063	[0.61664, 1.0921]	0.1752
Fri	1.0351	[0.78323, 1.3679]	0.80856	1.0164	[0.76888, 1.3436]	0.90906
Sat	0.63362	[0.41744, 0.96175]	**0.032101**	0.65366	[0.43083, 0.99175]	**0.045619**
Sun	1.1682	[0.52533, 2.598]	0.70296	1.1188	[0.5038, 2.4848]	0.78266
Month (reference: October)						
Nov	0.92421	[0.56991, 1.4987]	0.74931	0.88905	[0.549, 1.4397]	0.63255
Dec	0.75408	[0.46276, 1.2288]	0.25723	0.73494	[0.45054, 1.1989]	0.2174
Jan	1.2419	[0.79685, 1.9354]	0.33863	1.1316	[0.72329, 1.7703]	0.58827
Feb	1.2952	[0.82775, 2.0267]	0.25746	1.2044	[0.77025, 1.8832]	0.41485
Mar	1.1812	[0.75061, 1.8588]	0.47158	1.086	[0.69014, 1.7089]	0.72134
Apr	1.1433	[0.72673, 1.7987]	0.5624	1.0667	[0.6772, 1.6802]	0.78064
May	1.1386	[0.72432, 1.7899]	0.57376	1.1082	[0.70536, 1.7411]	0.65582
Jun	1.1323	[0.71137, 1.8022]	0.60041	1.0328	[0.6505, 1.6399]	0.89102
Jul	1.3544	[0.87266, 2.1019]	0.17618	1.2901	[0.83226, 1.9997]	0.25474
Aug	0.98789	[0.61995, 1.5742]	0.95911	0.95976	[0.6034, 1.5266]	0.86228
Sep	0.96974	[0.60957, 1.5427]	0.89679	0.93984	[0.59078, 1.4951]	0.79337

Table 12: Results of fitting the GLM for rationing, full day-of-week configuration.

	Block features			Rolling features		
Variable	Odds ratio	95% CI	p-value	Odds ratio	95% CI	p-value
Intercept and resource availability						
Intercept	0.93726	[0.56892, 1.5441]	0.79918	0.80186	[0.4874, 1.3192]	0.38466
THShare_{t-1}	1.0063	[0.75549, 1.3403]	0.96593	1.0967	[0.81044, 1.484]	0.5499
Holiday	0.77958	[0.39582, 1.5354]	0.4715	0.76008	[0.38439, 1.503]	0.43032
ES assignments	0.99617	[0.97692, 1.0158]	0.69997	1.0055	[0.98611, 1.0252]	0.5819
TH assignments	0.97677	[0.95351, 1.0006]	0.055911	0.98536	[0.96065, 1.0107]	0.25512
ES exits	1.0172	[0.99797, 1.0368]	0.079868	1.0264	[1.0072, 1.046]	**0.0068166**
per SD	1.0987	[0.98885, 1.2208]	0.079868	1.1604	[1.0419, 1.2925]	**0.0068166**
TH exits	0.98276	[0.95818, 1.008]	0.1783	0.95115	[0.92662, 0.97632]	**0.00017158**
per SD	0.92592	[0.82776, 1.0357]	0.1783	0.80218	[0.71506, 0.89992]	**0.00017158**
Day of the week (reference: Monday)						
Tue	0.94873	[0.72376, 1.2436]	0.70309	0.97212	[0.74027, 1.2766]	0.83881
Wed	0.8095	[0.61511, 1.0653]	0.13146	0.81628	[0.61899, 1.0765]	0.15042
Thu	0.60038	[0.44575, 0.80863]	**0.00078508**	0.61642	[0.45658, 0.83221]	**0.0015816**
Fri	0.77587	[0.58381, 1.0311]	0.080311	0.79535	[0.59735, 1.059]	0.11698
Sat	0.74236	[0.51675, 1.0665]	0.10702	0.76533	[0.53192, 1.1012]	0.14963
Sun	0.02654	[0.0036494, 0.19301]	**0.00033708**	0.026335	[0.0036196, 0.19161]	**0.00032845**
Month (reference: October)						
Nov	0.6548	[0.41116, 1.0428]	0.074522	0.66835	[0.41937, 1.0652]	0.090163
Dec	0.52942	[0.33328, 0.84099]	**0.0070731**	0.5479	[0.34428, 0.87196]	**0.011152**
Jan	0.55142	[0.35876, 0.84754]	**0.0066427**	0.60382	[0.3911, 0.93224]	**0.02281**
Feb	0.81319	[0.52885, 1.2504]	0.34621	0.82308	[0.53581, 1.2644]	0.37401
Mar	0.66739	[0.43831, 1.0162]	0.059432	0.69384	[0.45511, 1.0578]	0.089357
Apr	0.76344	[0.49625, 1.1745]	0.21939	0.81723	[0.53023, 1.2596]	0.3605
May	0.84075	[0.55147, 1.2818]	0.42013	0.85311	[0.55915, 1.3016]	0.4611
Jun	0.51989	[0.33735, 0.80121]	**0.0030332**	0.52079	[0.33747, 0.80371]	**0.0032082**
Jul	0.68562	[0.44465, 1.0572]	0.087569	0.71611	[0.4642, 1.1047]	0.13114
Aug	0.70711	[0.46152, 1.0834]	0.11136	0.71716	[0.46774, 1.0996]	0.12735
Sep	0.90026	[0.58582, 1.3835]	0.63173	0.9493	[0.61692, 1.4608]	0.81296

H.3 Remaining slack over the fiscal year

The model allows sensitivity to remaining budget to strengthen over part of the horizon, but not necessarily through its end: once budget is ample, the spending probability is $1 - p$ regardless of the positive budget left. We therefore treat a remaining-slack interaction with fiscal position as a descriptive check for strengthening capacity associations. We test this by interacting a measure of remaining slack with the position in the fiscal year. Remaining slack on day t is the cumulative number of transitional housing (TH) exits minus TH entries on earlier days of the same fiscal year (October 1 to September 30), minus its average for the same fiscal week across years, standardized to mean zero and standard deviation (SD) one. Fiscal position is the number of days since October 1 divided by 365. We add slack to the rolling day-type model of Table 1 (the slack-only model) and then add its product with fiscal position (the interaction model). The sample is the 1,329 days with ES or TH placements in fiscal years 2009–2013; fiscal year 2008 is excluded because the daily series begins partway through it. Because 26.1% of TH entries are upgrades, net slack partly reflects past overrides, so we also report a version built from TH exits alone.

Table 13: Slack terms of the remaining-slack test, fiscal years 2009 to 2013. Both models are the rolling day-type model of Table 1: the slack-only model adds slack, and the interaction model adds slack and the product of fiscal position and slack; the overall model counts overrides among placements predicted ES or TH. Slack is standardized, so its odds ratios are per SD, and the interaction is the factor by which the effect of slack changes over a full fiscal year. The slack term of the interaction model, the effect of slack on October 1, is not shown, because slack is close to 0 then by construction. p-values below 0.05 are in bold.

		Slack (slack-only model)			Fiscal position × slack (interaction model)		
Model	Days	Odds ratio	95% CI	p-value	Odds ratio	95% CI	p-value
Net slack (TH exits minus TH entries)							
Overall	1,322	0.953	[0.877, 1.035]	0.250	1.083	[0.826, 1.420]	0.562
Upgrading	1,191	0.879	[0.779, 0.991]	**0.036**	1.284	[0.862, 1.912]	0.219
Rationing	979	1.112	[0.983, 1.258]	0.092	0.806	[0.540, 1.202]	0.289
Exits only (TH exits)							
Overall	1,322	0.988	[0.913, 1.069]	0.757	1.013	[0.675, 1.520]	0.952
Upgrading	1,191	1.012	[0.905, 1.133]	0.830	0.775	[0.431, 1.393]	0.394
Rationing	979	0.908	[0.804, 1.025]	0.119	1.792	[0.971, 3.305]	0.062

Table 13 reports the results. The interaction is above one for overall discretion (1.083, $p = 0.562$) and upgrading (1.284, $p = 0.219$) and below one for rationing (0.806, $p = 0.289$), and no interaction is significant at the 5% level under either measure of slack. Adding slack also leaves the calendar pattern largely unchanged: the month coefficients of the overall model shrink by 3.7% in total, and no month odds ratio moves by more than 0.02. The one significant slack term, less upgrading with more net slack (0.879, $p = 0.036$), disappears with exits alone (1.012, $p = 0.830$), consistent with net slack partly reflecting past upgrades.

We find no evidence that the association between remaining slack and overrides strengthens over the fiscal year. These data therefore cannot separate budget rationing from calendar-based explanations, and we report the calendar patterns in Section 4 as descriptive rather than as evidence about the shape of the gain distribution.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The abstract and introduction (Section 1) clearly state the three main contributions and the main body of the paper (in conjunction with the appendix) expand on each of those contributions.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: The paper includes discussion of limitations in Section 5. All assumptions that the theory rests on are clearly pointed out throughout the paper.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer:[Yes]

Justification: Both main theorems (Theorem 3.1 and 3.2) state their assumptions explicitly. Proof sketches appear in the main text and complete proofs are provided in Appendix D and Appendix E. Supporting results are numbered, cross-referenced, and proved in the appendix.

Guidelines:

- The answer [N/A] means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: For the theoretical simulations, all distribution families, parameterizations, and sample sizes ($N = 10^6$) are reported in Appendix E.7. For the empirical analysis, the GLM specification, feature construction (rolling 7-day windows), aggregation procedure (daily binomial), and all covariates are described in Section 4 and Appendix G. The decision-tree baseline follows the SER-DT method as used by Pokharel et al. [33] with max depth 4 on 34 features (listed in Table 6). The HMIS dataset is administrative and not publicly releasable, but all methodology is fully specified.

Guidelines:

- The answer [N/A] means that the paper does not include experiments.
- If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).

(d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [No]

   Justification: The HMIS data is protected administrative data that we use under a specific data sharing agreement. It cannot be publicly released due to privacy restrictions. Simulation for the theoretical results (scale invariance, shape dependence) can be fully reproduced from the specifications in Appendix E.7.

   Guidelines:

   - The answer [N/A] means that paper does not include experiments requiring code.
   - Please see the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
   - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
   - The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
   - At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
   - Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [Yes]

   Justification: Simulation details (distributions, parameterizations, sample sizes) are in Appendix E.7. The empirical GLM specification is fully described in Appendix G.5, including the link function (logit), weighting scheme (trial counts), covariates, and reference categories. Feature engineering is detailed in Appendix G.4. Computing infrastructure is reported in the beginning of Appendix G.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: The empirical results in Table 1 report odds ratios with 95% confidence intervals and significance levels ($p < 0.001$, $p < 0.01$, $p < 0.05$, $p < 0.1$). Full coefficient tables with confidence intervals and p-values appear in Appendix H.

Guidelines:

- The answer [N/A] means that the paper does not include experiments.
- The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
- If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: Computing infrastructure is reported in the beginning of Appendix G.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: The research uses de-identified administrative HMIS records and does not involve direct interaction with human subjects. The work aims to improve understanding of discretionary decision-making in social services, with potential benefits for more equitable resource allocation.

   Guidelines:

   - The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.

- If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [Yes]

    Justification: Section 5 discusses implications for decision-support system design, noting that systems should help manage not only which cases merit discretion but also when to deploy it. The paper frames discretion as potentially welfare-improving when treated as an explicitly budgeted resource.

    Guidelines:

    - The answer [N/A] means that there is no societal impact of the work performed.
    - If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
    - Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
    - The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
    - The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
    - If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

    Answer: [N/A]

    Justification: The paper does not release models, datasets, or tools with high risk for misuse. The contribution is a theoretical framework and empirical analysis of existing administrative data.

    Guidelines:

    - The answer [N/A] means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [Yes]

    Justification: The HMIS data is cited via relevant prior work [23, 33]. We also cite the code we use: the SER-DT reference implementation [39] (`https://github.com/victorfcsouza/explainable-dts`, commit 86d017e, run unchanged; the repository does not state a license), GOSDT-guesses [27] (`gosdt` 1.0.4, BSD 3-Clause license), and XGBoost [7] (version 3.4.1, Apache 2.0 license). No pre-trained models are used.

    Guidelines:

    - The answer [N/A] means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
    - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
    - If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
    - For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
    - If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [N/A]

    Justification: The paper does not release new datasets, models, or code packages.

    Guidelines:

    - The answer [N/A] means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A]

    Justification: The paper does not involve crowdsourcing or direct research with human subjects.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.

- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer:[N/A]

    Justification: The study analyzes pre-existing de-identified administrative data and does not constitute human subjects research requiring IRB review. No interventions were performed on individuals.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
    - We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
    - For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

    Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

    Answer: [N/A]

    Justification: LLMs were not used as a component of the core methodology, theoretical development, or empirical analysis in this research.

    Guidelines:

    - The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
    - Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.