Sequential Offering in On-Demand Platforms: On the Optimality of Greedy Ranking

Hongyao Ma
Graduate School of Business, Columbia University, New York, NY 10027, hongyao.ma@columbia.edu
Will Ma
Graduate School of Business, Columbia University, New York, NY 10027, wm2428@gsb.columbia.edu
Matias Romero
Graduate School of Business, Columbia University, New York, NY 10027, mer2262@columbia.edu

On-demand platforms face the fundamental challenge of fulfilling time-sensitive jobs with independent workers who may decline offers. To minimize delays and unfulfilled jobs, platforms frequently raise the offered wage sequentially following each rejection. However, the interaction between these dynamic price adjustments and the specific sequence in which workers are approached has been overlooked. In particular, if the best-suited workers (e.g., closest to the job) are also ranked earliest in the sequence, then those workers would see the lowest offered wages and may decline, leading to poor system outcomes where less-suited workers end up seeing the raised wages and accepting the job.

We study the sequential offering problem to maximize expected welfare or platform profit by jointly optimizing the ranking of workers and the pricing trajectory. Surprisingly, our main result establishes that if the reservation wage distribution exhibits a non-increasing and convex density function (e.g., Uniform, Exponential, etc), the exact welfare-optimal solution is achieved by greedy ranking and dynamically optimizing wages via backward induction. Numerical experiments further demonstrate that greedy ranking remains optimal or near-optimal across a broad range of other distributions. This suggests that rather than sending initial "low ball" offers to worse matches, platforms should stick with greedy ranking and optimize the wage offerings by appropriately taking the continuation value of the downstream offers into consideration.

1. Introduction

On-demand platforms face the challenge of fulfilling time-sensitive jobs (i.e., customer orders or requests) without the authority to enforce acceptance from workers (e.g. drivers, couriers, etc). Instead, a job is offered to workers sequentially, and is often rejected multiple times before a worker accepts the job. These rejections are largely driven by workers' sensitivity to offered wages (Allon et al. 2023). With job characteristics and wage now shown upfront across major platforms, the resulting cherry-picking has led to widespread reports of single-digit acceptance rates (Bloomberg 2021, Reddit 2023).

To reduce delays and unfulfilled jobs, both of which entail significant costs, platforms often increase the offered wage after each rejection. This strategy has been widely adopted across diverse domains, spanning food delivery (DoorDash 2025), ride-sharing (Uber 2024), freight marketplaces (Amazon Flex 2023), and healthcare staffing (CareRev 2023). While such dynamic wage adjustments effectively improve reliability, the interaction with the sequence (i.e., ranking) in which workers are chosen has been largely overlooked. Specifically, greedily ranking the best worker first can be inefficient when the offered wages first "low ball" the worker who are good matches, and later incentivize the poor matches to accept the job with higher wages.

To illustrate, consider a simple scenario of offering a job between two workers, where one is a better match than the other. Suppose the platform makes an initial "low ball" offer of \$5 that is always rejected. If rejected, the wage increases for the second worker to \$10, that is accepted by most workers. Under a greedy ranking, the best-matching worker rejects the first offer and the worse-matching worker fulfills the job.

Recognizing how dynamic wages alter the likelihood of acceptance, platforms are increasingly exploring how to adjust their initial rankings in response to these anticipated wage sequences. However, attempting to solve this joint optimization problem iteratively creates a complex feedback loop that is not guaranteed to resolve the underlying inefficiency.

In this paper, we propose the following simple method for jointly optimizing the ranking and wages: fix the ranking to be Greedy, and then optimize the wages knowing this would be the case, which would avoid the issue above. This reduces the intractable search problem over all permutations of workers to a single sorting problem, where the workers are ordered in decreasing quality of match with the job. The optimal wages given this order can then be tractably found using dynamic programming.

We show that surprisingly, this simple Greedily-rank-then-optimize-wages method is actually optimal for many distributions of worker willingness to work. We now elaborate on our model and results..

1.1. Main Contributions

We study a sequential offering problem where a platform must jointly optimize the sequence in which workers are approached (i.e. ranking) and the dynamic wage trajectory offered to them. Although our primary analytical results focus on welfare optimization, we demonstrate that the platform still strategically utilizes lower initial wages as a mechanism of self-selection, aiming for the job to be ultimately fulfilled by the better match. We also show in Appendix B that the optimality of greedy ranking can be extended to profit maximization.

Exact optimality of greedy ranking. We establish in Theorem 1 that if the distribution F admits a non-increasing and convex density function, the exact welfare-optimal solution is achieved by greedy ranking and dynamically optimizing wages via backward induction. These conditions correspond to settings where wage incentives yield diminishing marginal returns on acceptance probabilities at a decelerating rate, a property satisfied by many standard distributions (Uniform, Exponential, Pareto, etc).

Empirical robustness. Beyond our theoretical sufficient conditions, extensive numerical simulations demonstrate strong empirical evidence that the greedy ranking sequence remains optimal across a broad class of distributions (**??**). Motivated by these empirical findings, we develop a computer-assisted proof to formally certify the exact global optimality of the greedy sequence for several standard distributions. Even in environments where we force the greedy heuristic to fail (e.g., highly skewed Weibull distributions), the observed optimality gap never exceeds 10% in the worst case, and drops to less than 1% on average.

Worst-case guarantees. To complement our empirical observations and formalize the performance limits of the greedy heuristic when it is suboptimal, we establish tight worst-case guarantees. We know from the prophet inequality literature and its connection to sequential pricing (Correa

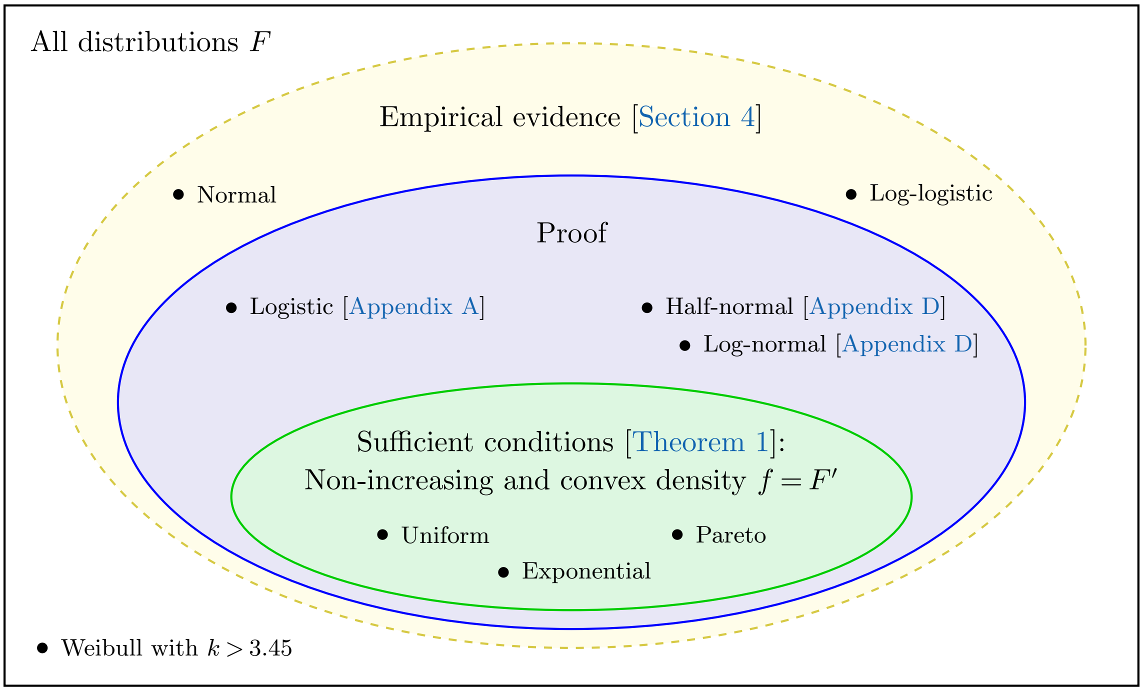

Figure 1 Classification of probability distributions F based on the level of evidence supporting the exact global optimality of the greedy ranking sequence.

et al. 2019) that $1/2$ serves as a universal lower bound relative to the offline optimum. Although this competitive ratio remains asymptotically tight in our setting, we demonstrate that our additional structure allows for a strict improvement over $1/2$ for any finite n: the greedy policy guarantees at least a $n/(2n-1)$ fraction of the prophet benchmark.

1.2. Connection to an open theoretical problem in prophet inequalities

Our joint optimization problem can be framed as a *free-order* prophet inequality problem (see Hill 1983, for problem definition). In the general formulation, determining the optimal sequencing rule is computationally intractable and known to be NP-hard (Agrawal et al. 2020). However, our environment exhibits a specific structure: rather than the general problem, we are studying an IID special case with location shifts, where each worker's acceptance threshold is determined by a heterogeneous parameter.

Because of this structure, we ask a very different question, which is *characterizing the set of distributions* for which the greedy-rank-then-optimize-wages heuristic is *optimal* under *any* location shifts. A similar question has in fact previously intrigued **?**; however, to the best of our knowledge, our paper represents the only progress on this type of question since then.

1.3. Further Related Work

Matching and pricing in on-demand platforms. Our work contributes to the extensive literature on pricing and matching in on-demand platforms. In this domain, the initial focus is on demand-side pricing and spatial-temporal matching to balance overall network capacity, such as using surge pricing to manage customer demand and network efficiency (Castillo et al. 2017, Bimpikis et al. 2019). On the supply side, existing research primarily evaluates optimal wage contracts and commission structures under the assumption of aggregate supply responses or self-scheduling capacity (Cachon et al. 2017, Taylor 2018). Building on this, recent literature has increasingly examined driver pricing and scheduling through a macroscopic, equilibrium lens (Besbes et al. 2021, Ma et al. 2022, Garg and Nazerzadeh 2022). These approaches generally assume the platform acts as a macroscopic market maker, optimizing market-clearing prices across a broader pool of agents rather than executing sequential, targeted offers to specific heterogeneous workers.

Sequencing and sequential pricing. The stochastic probing literature largely focuses on algorithmic approximations for environments with exogenous acceptance probabilities. Purohit et al. (2019) formalize the hiring under uncertainty problem, demonstrating that greedy heuristics are generally suboptimal under arbitrary acceptance probabilities. Epstein and Ma (2024) employ linear programming relaxations and dependent rounding to establish constant-factor approximations for sequential and parallel offering models subject to strict capacity constraints. A fundamental distinction in our setting is that acceptance probabilities are endogenous via the optimal wage decision. Surprisingly, this can make the optimal policy easier to compute by preventing adversarial sequence of acceptance probabilities.

A second stream of literature scales sequential decision-making to endogenous settings, where platform actions jointly dictate acceptance probabilities and rewards. Performance in this domain is typically evaluated against offline prophet benchmarks. While allowing the decision-maker to optimize the sequence—the *free-order* variant—improves competitive ratios (Beyhaghi et al. 2021, Peng and Tang 2022), the analytical focus remains reliant on fractional approximations. Extending these frameworks to two-sided networks, Pollner et al. (2024) analyze sequential pricing on bipartite graphs. More recently, Derakhshan et al. (2026) investigate an action-reward framework, utilizing a configuration linear program and an efficient polynomial-time approximation scheme (EPTAS) oracle to bypass severe integrality gaps. Our work departs from this reliance on algorithmic approximation; we return to exact structural optimality, mapping the specific distributional conditions under which exact dynamic programming aligns perfectly with simple greedy sequencing.

2. The Model

We consider the sequential offering problem faced by on-demand platforms. The platform seeks to fulfill a job, which yields a value v if completed, by matching it with one worker from a finite pool $I = [n]$.

To fulfill the job, worker $i \in I$ incurs a total cost of $c_i + Z_i$, where c_i is a deterministic scalar that captures platform-observable components, and Z_i is a random variable with cumulative distribution function (CDF) F representing private factors. For instance, in food delivery platforms, c_i can account for attributes such as heterogeneous locations and expected travel times, as well as batching efficiencies if the worker already has another job in their bag. Within this same context, Z_i captures unobservable idiosyncratic preferences that are treated as independent and identically distributed (i.i.d.) across workers to comply with fairness norms and regulations against behavioral price discrimination. Given this cost structure, worker i will accept a *total* wage of w_i^{total} if and only if it compensates for their total cost, $w_i^{\mathsf{total}} \geq c_i + Z_i$. Upon acceptance, the platform realizes a profit of $v - w_i^{\mathsf{total}}$, and the worker derives a surplus of $w_i^{\mathsf{total}} - c_i - Z_i$. The total welfare generated by the match is the sum of these payoffs, $v - c_i - Z_i$.

We normalize the platform's offer and the job's value against the observable cost c_i by defining the *net wage* as $w_i \triangleq w_i^{\mathsf{total}} - c_i$ and the *net value* as $v_i \triangleq v - c_i$. Under these definitions, the worker accepts the offer when their private cost falls below the net wage ($Z_i \leq w_i$), yielding an acceptance probability of $F(w_i)$. If they accept, the platform's profit is $v_i - w_i$ and the total welfare generated is $v_i - Z_i$. Thus, the problem environment is fully characterized by the distribution F and net values $\boldsymbol{v} = (v_1, \ldots, v_n) \in \mathbb{R}^n$.

The platform's problem consists of making sequential take-it-or-leave-it offers, jointly deciding the sequence in which workers are approached and the net wage to offer. Formally, the platform selects a *ranking* policy σ, defined as a bijection $\sigma : [n] \to I$, that dictates the order in which workers are approached, alongside a sequence of net wages $\boldsymbol{w} = (w_1, w_2, \ldots, w_n) \in \mathbb{R}^n$. At each step $t = 1, \ldots, n$, the platform approaches worker $\sigma(t)$ and offers w_t. With probability $F(w_t)$, the offer is accepted, and the job is fulfilled, yielding welfare of $v_{\sigma(t)} - Z_{\sigma(t)}$.[1] If the offer is rejected, the platform proceeds to step $t + 1$. Since private costs $(Z_i)_i$ are i.i.d., the platform's objective depends on the ranking policy solely through the induced sequence of net values. Let $\Pi(\boldsymbol{v}) \triangleq \{(v_{\sigma(1)}, v_{\sigma(2)}, \ldots, v_{\sigma(n)}) \mid \sigma : [n] \to I \text{ is bijective}\}$ denote the set of all such permutations. We streamline our notation by evaluating policies directly via their corresponding sequence $\boldsymbol{u} \in \Pi(\boldsymbol{v})$, expressing the expected welfare under net wages $\boldsymbol{w}$ as

$$V(\boldsymbol{u}, \boldsymbol{w}) = \sum_{t=1}^{n} \left(\prod_{s=1}^{t-1} (1 - F(w_s)) \right) F(w_t)\, \mathbb{E}[u_t - Z \mid Z \le w_t]. \tag{1}$$

To evaluate and compare ranking policies under their respective optimal wages, we overload the notation and define the maximum expected welfare achievable under sequence $\boldsymbol{u}$ as $V(\boldsymbol{u})$, and the performance of the optimal policy as $\mathsf{OPT}(\boldsymbol{v}) \triangleq \max_{u \in \Pi(\boldsymbol{v})} V(\boldsymbol{u})$. We assume without loss of generality that the worker pool is indexed such that $v_1 \ge v_2 \ge \cdots \ge v_n$, such that the *greedy ranking policy* is represented by $\boldsymbol{v} \in \Pi(\boldsymbol{v})$, and we denote its performance by $\mathsf{GRE}(\boldsymbol{v}) \triangleq V(\boldsymbol{v})$.

This sequential offering problem fits into the broad literature on sequential posted pricing and optimal stopping problems. In particular, for any given sequence, the optimal wages can be solved efficiently via dynamic programming. The sequencing decision, however, poses a significant theoretical challenge: although irrelevant in a purely i.i.d. environment, it is computationally NP-hard

[1] While our main exposition focuses on welfare, in Appendix B we generalize our framework to capture profit maximization.

when arbitrarily heterogeneous distributions are allowed (Agrawal et al. 2020). Crucially, our setting exhibits a *location-shift structure*. Because workers are homogeneous in their distribution of unobservable private costs, their reservation wage distributions are simply shifted by net values v_i. This specific structure interpolates between the two contrasting regimes and motivates a natural heuristic: a greedy ranking policy that approaches workers in descending order of their observable net values.

In what follows, we leverage this structure to characterize the optimal wages and, more importantly, to derive a tractable recursion for $V(\boldsymbol{u})$ that allows us to systematically analyze and compare different ranking policies.

3. Optimality of Greedy Ranking

In this section, we establish the exact optimality of the greedy ranking policy. We first derive a closed-form recursive solution for the optimal wages and welfare objective under an arbitrary sequence, and then leverage the structure to study adjacent pairwise swaps.

3.1. Optimal Wages

Given a sequence $\boldsymbol{u} \in \Pi(\boldsymbol{v})$, the wage optimization problem is a finite-horizon dynamic program over the sequence of workers. Let $V(\boldsymbol{u}_{t:n})$ denote the optimal expected continuation value derived from the remaining sequence of workers from step t to n, with the convention that $V(\boldsymbol{u}_{n+1:n}) = 0$. At each step $t \in [n]$, the platform decides the wage w_t to offer the current worker, balancing the immediate expected welfare against the continuation value. The Bellman equation is given by:

$$\begin{aligned} V(\boldsymbol{u}_{t:n}) &= \max_{w_t \in \mathbb{R}} \left\{ F(w_t)\big(u_t - \mathbb{E}[Z \mid Z \le w_t]\big) + \big(1 - F(w_t)\big) V(\boldsymbol{u}_{t+1:n}) \right\} \\ &= V(\boldsymbol{u}_{t+1:n}) + \max_{w_t \in \mathbb{R}} \left\{ F(w_t)\big(u_t - \mathbb{E}[Z \mid Z \le w_t] - V(\boldsymbol{u}_{t+1:n})\big) \right\}. \end{aligned} \tag{2}$$

Notice that the optimal wage depends entirely on the marginal value of the current worker on top of the downstream continuation value, $x \triangleq u_t - V(\boldsymbol{u}_{t+1:n})$. To capture the expected incremental welfare generated at any given step, we define the *incremental value function*

$$\psi(x) \triangleq \max_{w \in \mathbb{R}} \left\{ F(w)\big(x - \mathbb{E}[Z \mid Z \le w]\big) \right\}. \tag{3}$$

This function isolates the difference between the expected welfare objective at two consecutive steps. The following lemma establishes that it is welfare-optimal to set the net wage equal to the marginal value, causing worker $\sigma(t)$ to accept whenever $V(\boldsymbol{u}_{t+1:n}) \leq u_t - Z$, i.e. whenever the welfare from them accepting exceeds the value-to-go $V(\boldsymbol{u}_{t+1:n})$.

Lemma 1. *For any marginal value $x \in \mathbb{R}$, the maximization in the incremental value function is solved by setting $w = x$, which evaluates to:*

$$\psi(x) = \mathbb{E}[\max\{0, x - Z\}] = \int_{-\infty}^{x} F(s)ds. \tag{4}$$

Consequently, the optimal wage at step t is $w_t = u_t - V(\boldsymbol{u}_{t+1:n})$, and the optimal continuation values satisfy the recursion:

$$V(\boldsymbol{u}_{t:n}) = V(\boldsymbol{u}_{t+1:n}) + \psi\big(u_t - V(\boldsymbol{u}_{t+1:n})\big), \quad t \in [n]. \tag{5}$$

Proof. Let $x \in \mathbb{R}$. Notice that for any wage $w \in \mathbb{R}$, we have $F(w)\mathbb{E}[Z \mid Z \leq w] = \mathbb{E}[Z\mathbb{1}\{Z \leq w\}]$. Hence, the maximization objective in Equation (3) simplifies to $\max_{w\in\mathbb{R}} \mathbb{E}[(x - Z)\mathbb{1}\{Z \leq w\}]$, which is at most $\mathbb{E}[\max\{0, x - Z\}]$. This upper bound is achieved by setting $w = x$, which aligns the indicator to include all and only the realizations where $x - Z \geq 0$. Finally, by expressing $\max\{0, x - Z\} = \int_{-\infty}^{x} \mathbb{1}\{Z \leq s\}ds$, we can apply Tonelli's Theorem to interchange the expectation and integral, completing the proof. □

3.2. Local Stability

Before proving global optimality, we establish that when wages are correctly optimized for the greedy sequence $\boldsymbol{v} \in \Pi(\boldsymbol{v})$, no other ranking achieves higher expected welfare at those specific wages. This ensures that if the current wages are optimized for the greedy sequence, then updating the sequence alone cannot improve the welfare. This result acts as an essential component of our global optimality proof.

To formalize this, consider two adjacent workers in a greedy sequence at positions t and $t+1$. For the greedy sequence to be *stable* against a pairwise swap, the unconditional probability of the downstream worker accepting the offer cannot strictly exceed the unconditional probability of the upstream worker accepting. If it did, the platform could strictly improve expected welfare by swapping the two workers while holding the offered wages fixed, since the upstream worker has a weakly higher observable value ($v_t \geq v_{t+1}$). Therefore, stability requires:

$$F(w_{t+1})\big(1 - F(w_t)\big) \leq F(w_t), \quad \forall t \in [n]. \tag{6}$$

Let $x = v_{t+1} - V(\boldsymbol{v}_{t+2:n})$ and $y = v_t - V(\boldsymbol{v}_{t+2:n})$ denote their respective marginal values over the shared future continuation value. By Lemma 1, it is welfare-optimal to set $w_{t+1} = x$ and $w_t = y - \psi(x)$. Substituting these into the stability condition and rearranging gives $F(y - \psi(x)) \geq \frac{F(x)}{1+F(x)}$. Because $y \geq x$ and F is non-decreasing, the left-hand side is naturally bounded below by $F(x - \psi(x))$. Thus, to guarantee local stability for all possible value pairs, it is sufficient that the distribution satisfies:

$$F(x - \psi(x)) \geq \frac{F(x)}{1 + F(x)}, \quad \forall x \in \mathbb{R}. \tag{7}$$

The following lemma establishes that this property is satisfied when the cost distribution F is concave, representing diminishing marginal returns in acceptance probability.

LEMMA 2. *Assume the cumulative distribution function F is continuous, satisfies $F(0) = 0$, and has a non-increasing density function. Then, Equation (7) holds and greedy ranking is stable under welfare-optimal wages.*

Proof. Using the linearity of integrals, we rewrite the argument as $x - \psi(x) = \int_0^x (1 - F(t))dt$. This integral exactly represents the expected value of $\min(Z, x)$, where $Z \sim F$. Let $Z_x = \min(Z, x)$. Since $Z \geq 0$, the expected value is $\mathbb{E}[Z_x] = x - \psi(x)$.

Because the density function is non-increasing, F is concave on its support. By Jensen's Inequality, $F(\mathbb{E}[Z_x]) \geq \mathbb{E}[F(Z_x)]$. Evaluating the right-hand side yields $\mathbb{E}[F(Z_x)] = \int_0^x F(t)dF(t) + \mathbb{P}(Z > x)F(x) = \frac{1}{2}F(x)^2 + (1 - F(x))F(x) = F(x) - \frac{1}{2}F(x)^2$.

Therefore, Jensen's inequality implies $F(x - \psi(x)) \geq F(x) - \frac{1}{2}F(x)^2$. To verify that $F(x) - \frac{1}{2}F(x)^2 \geq \frac{F(x)}{1+F(x)}$, multiply both sides by $1 + F(x) \geq 0$, yielding $F(x) + \frac{1}{2}F(x)^2(1 - F(x)) \geq F(x)$. Since $F(x) \in [0, 1]$, this inequality is always satisfied, completing the proof. $\square$

3.3. Global Optimality

In the previous subsection, we established that the greedy ranking is locally stable when paired with its optimal wages, provided the cost distribution features diminishing marginal returns. However, local stability between two workers does not automatically guarantee that the greedy assignment is globally optimal against all possible permutations.

To prove global optimality, we use a swapping argument. Consider an arbitrary non-greedy sequence $\boldsymbol{u} \in \Pi(\boldsymbol{v}) \setminus \{\boldsymbol{v}\}$ and identify an adjacent inverted pair, where the worker at the upstream position has a strictly lower observable value than the worker at the downstream position. We then evaluate the impact of swapping their positions to place the superior worker first. Crucially, simply flipping the workers' positions while holding their offered wages constant is insufficient and will generally decrease welfare. Because the optimal wage is intrinsically linked to the continuation value, altering the order fundamentally changes the mathematically correct wage for the new upstream position. Therefore, we must dynamically recalculate the optimal prices for the newly swapped pair.

Let t and $t+1$ be the positions of an adjacent inverted pair, such that $u_t < u_{t+1}$. We define the marginal values of these two workers over the shared future continuation value as $x = u_t - V(\boldsymbol{u}_{t+2:n})$ and $y = u_{t+1} - V(\boldsymbol{u}_{t+2:n})$. Because the pair is inverted, we have $x < y$. Applying the recursion from Lemma 1, we directly evaluate the expected continuation value extracted from this pair in the original non-greedy order:

$$
\begin{aligned}
V(\boldsymbol{u}_{t+1:n}) &= V(\boldsymbol{u}_{t+2:n}) + \psi(u_{t+1} - V(\boldsymbol{u}_{t+2:n})) \\
&= V(\boldsymbol{u}_{t+2:n}) + \psi(y), \text{ and} \\
V(\boldsymbol{u}_{t:n}) &= V(\boldsymbol{u}_{t+1:n}) + \psi\big(u_t - V(\boldsymbol{u}_{t+1:n})\big) \\
&= V(\boldsymbol{u}_{t+2:n}) + \psi(y) + \psi\big(x - \psi(y)\big).
\end{aligned}
$$

If we swap their positions and re-optimize the wages, the downstream worker now has margin x, adding $\psi(x)$ to the continuation value $V(\boldsymbol{u}_{t+2:n})$. The upstream worker faces a margin of $y - \psi(x)$, resulting in a new expected continuation value of $V(\boldsymbol{u}_{t+2:n}) + \psi(x) + \psi(y - \psi(x))$.

We define the swap difference $D(x,y)$ as the expected welfare difference between the swapped sequence and the original sequence:

$$D(x,y) \triangleq \psi\big(y-\psi(x)\big)+\psi(x)-\psi\big(x-\psi(y)\big)-\psi(y). \tag{8}$$

If $D(x,y)\geq 0$ for all $x\leq y$, then this swap weakly improves expected welfare. By iteratively swapping adjacent inverted pairs, any non-greedy sequence can be transformed into the greedy sequence $\boldsymbol{v}$ without decreasing expected welfare. This establishes that no sequence strictly outperforms the greedy sequence.

To guarantee that this swap difference is always non-negative, we impose an additional assumption on the distribution. While local stability in Lemma 2 is satisfied by a non-increasing density, global optimality is guaranteed if we further assume this density is convex.

THEOREM 1. *Assume the cumulative distribution function F is continuous, satisfies $F(0)=0$, and admits a convex probability density function f, then* $\mathsf{GRE}(\boldsymbol{v})=\mathsf{OPT}(\boldsymbol{v})$ *for all* $\boldsymbol{v}\in\mathbb{R}^n$.

Proof. Fix the superior worker's margin $y\geq 0$ and consider the swap difference $D(x,y)$ as a function of the inferior worker's margin x on the domain $x\leq y$. At the boundary where the workers are identical, substituting $x=y$ into Equation (8) yields $D(y,y)=0$.

We evaluate the first partial derivative with respect to x:

$$\frac{\partial D}{\partial x}(x,y)=\psi'(x)\big[1-\psi'\big(y-\psi(x)\big)\big]-\psi'\big(x-\psi(y)\big). \tag{9}$$

Evaluating this derivative at the boundary $x=y$ yields $\frac{\partial D}{\partial x}(y,y)=\psi'(y)[1-\psi'(y-\psi(y))]-\psi'(y-\psi(y))$. Because $\psi'(w)=F(w)$, this recovers the local stability condition established in Equation (7). Since F is concave, $\frac{\partial D}{\partial x}(y,y)\leq 0$.

To prove that $D(x,y)$ never drops below zero for $x<y$, it is sufficient to show that its partial derivative $\frac{\partial D}{\partial x}(x,y)$ is a convex function of x. If this derivative is convex and terminates at a non-positive value at $x=y$, it crosses the horizontal axis at most once from positive to negative. This property forces $D(x,y)$ to be quasi-concave, ensuring it remains non-negative on the domain $x\leq y$.

We verify this convexity by evaluating the third partial derivative. Taking two additional derivatives of $\frac{\partial D}{\partial x}(x, y)$ with respect to x, and substituting $\psi' = F$, $\psi'' = f$, and $\psi''' = f'$, yields:

$$\frac{\partial^3 D}{\partial x^3}(x,y) = f'(x)\big[1 - F\big(y - \psi(x)\big)\big] + 3F(x)f(x)f\big(y - \psi(x)\big) - F(x)^3 f'\big(y - \psi(x)\big) - f'\big(x - \psi(y)\big). \tag{10}$$

We group the terms of this third derivative to show that the expression is non-negative:

1. The term $3F(x)f(x)f(y - \psi(x))$ is non-negative because both the cumulative distribution and the probability density are non-negative.
2. The term $-F(x)^3 f'(y - \psi(x))$ is non-negative because the concavity of F requires its density f to be non-increasing, implying $f' \leq 0$.
3. We group the remaining terms into $[f'(x) - f'(x - \psi(y))] - f'(x)F(y - \psi(x))$. Because f is convex, its derivative f' is non-decreasing. Since $\psi(y) \geq 0$, we have $x - \psi(y) \leq x$, which implies $f'(x) - f'(x - \psi(y)) \geq 0$. The remaining component $-f'(x)F(y - \psi(x))$ is non-negative since $f' \leq 0$ and $F \geq 0$.

Summing these components confirms that $\frac{\partial^3 D}{\partial x^3}(x, y) \geq 0$. Consequently, the first partial derivative $\frac{\partial D}{\partial x}(x, y)$ is convex in x, forcing $D(x, y)$ to be quasi-concave. Since $D(y, y) = 0$, it follows that $D(x, y) \geq 0$ for all $x \leq y$. Therefore, any adjacent swap toward the greedy sequence weakly improves expected welfare, completing the proof. □

Theorem 1 establishes that the greedy sequence is globally optimal when the cost distribution exhibits a non-increasing, convex density. This structural condition corresponds to environments where wage incentives yield diminishing marginal returns on acceptance probabilities at a decelerating rate. While these properties are satisfied by several standard distributions, they serve as sufficient rather than strictly necessary conditions. In the following section, we explore the empirical robustness of the greedy heuristic when these market conditions are not satisfied.

4. Numerical Experiments

In Section 3, we established that greedy ranking is welfare-optimal when the private cost distribution exhibits a non-increasing and convex probability density function. These conditions correspond to settings where wage incentives yield diminishing marginal returns on acceptance probabilities at a decelerating rate, a property satisfied by many standard distributions (Uniform, Exponential, Pareto, etc). In this section, we study the empirical performance of the greedy ranking policy across a broad class of distributions that do not satisfy our sufficient conditions from Theorem 1.

4.1. Simulation Setup

For every tested distribution, we first search for net values that make the greedy ranking suboptimal. As established in Section 3, it is sufficient to search environments with $n = 2$ workers. We search over net values between 0 and 5 times the mean of the distribution, using steps of 0.01. If we find net values for which the greedy policy is suboptimal, we then explore how the suboptimality gap changes for larger sequence lengths ($n > 2$).

4.2. Numerical Success and Further Formal Certification

Our numerical evaluation reveals widespread numerical success for most tested distributions, including the Normal, Log-normal, Logistic, and Log-logistic distributions. Indeed, under these distributions, our simulations do not find any pair of net values for which the greedy sequence is suboptimal. To take these results a step beyond partial empirical evidence, we sought formal certification for these environments. In Appendix D, we developed a rigorous computer-assisted proof framework that transforms the continuous domain into a finite, adaptive grid, allowing us to formally certify the exact global optimality of the greedy ranking for the Half-Normal and Log-normal distributions. Furthermore, for the Logistic case, exact optimality can be proven analytically (see Appendix A).

4.3. Suboptimal Environments

Finding standard continuous distributions under which greedy ranking becomes suboptimal is surprisingly difficult. One systematic way to identify such instances is to break the local stability condition defined in Subsection 3.2. Specifically, taking the marginal value $x \to \infty$ in Equation (7)

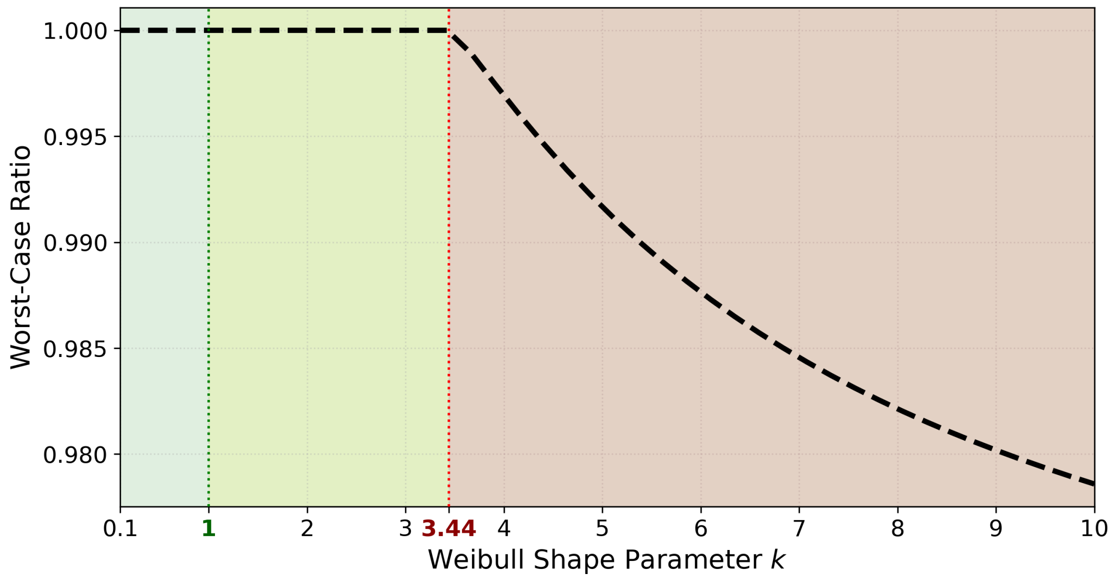

Figure 2 Worst-case optimality ratio of the greedy heuristic under a Weibull distribution ($\lambda = 1.0$) for $n = 2$.

reveals a fundamental boundary for the optimality of greedy ranking. Since $F(x) \to 1$ and $x - \psi(x) = \mathbb{E}[\min(Z, x)] \to \mathbb{E}[Z]$, a necessary condition for the greedy ranking to remain stable at large margins is that the mean of the base distribution must be weakly greater than its median. This motivates the study of a Weibull distribution with shape parameter k and scale λ, since its mean is defined by the Gamma function $\lambda\Gamma(1 + 1/k)$ and its median is $\lambda(\ln 2)^{1/k}$, meaning that the stability condition is violated whenever the shape parameter exceeds $k_0 \approx 3.44$.

We fix the scale parameter at $\lambda = 1.0$ and vary the shape parameter k to control the skewness of the distribution. Figure 2 shows the worst-case ratio, defined as the lowest observed ratio of greedy expected welfare to the dynamic optimum across our tested net values. For $k \leq 1$, the corresponding probability density function is decreasing and convex, satisfying the conditions of Theorem 1 that ensure the optimality of greedy ranking. In the intermediate regime $1 < k \leq k_0$, our numerical evaluation finds no suboptimal pairs of net values, providing empirical evidence that greedy ranking remains optimal. In contrast, for all parameters $k \geq k_0$, the mean lies below the median, and greedy ranking becomes suboptimal, with the ratio gradually decreasing as k grows large. Nonetheless, even when extending our numerical evaluation to extreme shape parameters up to $k = 100$, the minimum optimality ratio remains above 96%.

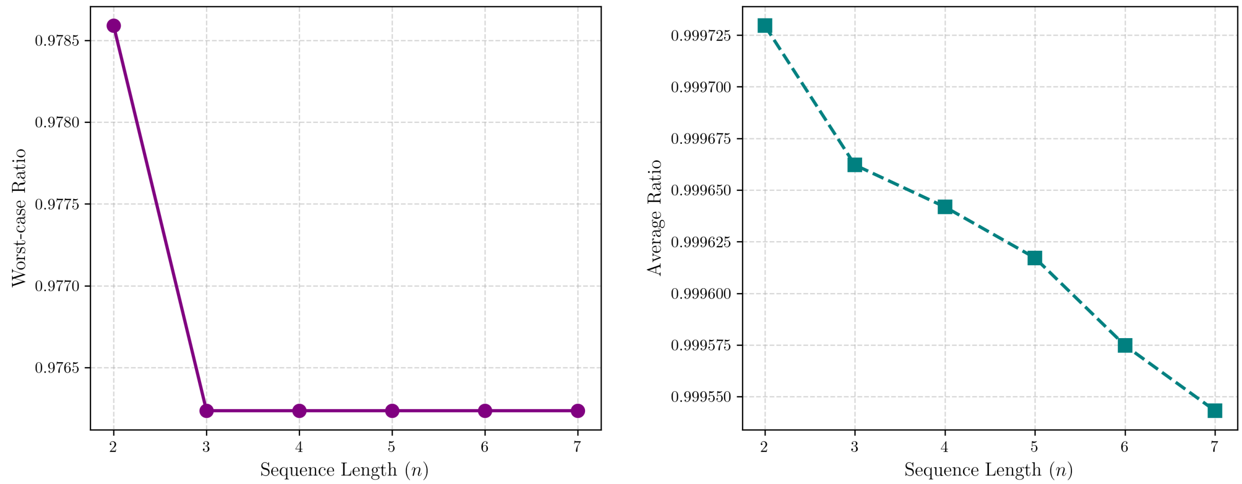

(a) Worst-case ratio. (b) Average ratio.

Figure 3 Performance of the greedy heuristic under a Weibull distribution ($k = 10$) as n increases.

To explore how this optimality gap evolves with the number of workers n, we fix the shape parameter at $k = 10$ and evaluate the all possible ranking policies across the same grid of net values. Figure 3a shows that the worst-case ratio initially decreases by roughly 0.25% when adding one extra worker at $n = 3$, but stabilizes for larger values of n. In contrast, the average ratio across the evaluated grid decays almost linearly as n increases, but remains above 99.9% (see Figure 3b).

These numerical results motivate the study of universal worst-case guarantees, which we establish in Section 5.

5. Worst-case Guarantees

We benchmark the greedy ranking against two baselines: the online optimum $\mathsf{OPT}(\boldsymbol{v}) = \max_{\boldsymbol{u}\in\Pi(\boldsymbol{v})} V(\boldsymbol{u})$, which optimizes over all sequences, and the offline prophet $\mathsf{OFF}(\boldsymbol{v}) \triangleq \mathbb{E}[\max_{i\in I}\{(v_i - Z_i)^+\}]$, which observes all cost realizations in hindsight. The classical prophet inequality establishes that even if we cannot choose the order, there is a universal lower bound of 1/2: for any distribution F, net values $\boldsymbol{v}$, and sequence length n, any online policy achieves at least half the prophet benchmark (Krengel and Sucheston 1977, Samuel-Cahn 1984).

This bound is tight in the general heterogeneous setting, where the classic worst-case instance constructs variables with vastly different variances. However, this instance cannot be cast under our location-shift structure, $X_i = (v_i - Z_i)^+$ with $Z_i \sim F$, because the underlying uncertainty is identical across all workers, locking the variance and tail shapes. Since the greedy ranking evaluates workers in decreasing order of observable values ($v_1 \geq \cdots \geq v_n$), the location-shift structure guarantees that the expected positive surplus from the current worker bounds that of any subsequent worker. By exploiting this monotonicity, we prove a tighter lower bound for finite sequences. Specifically, our bound is parameterized by n and performs better for small n, even though it also converges to the standard $1/2$ in the worst case as $n \to \infty$.

THEOREM 2. *For any sequence of $n \geq 1$ workers with net values $v_1 \geq v_2 \geq \cdots \geq v_n > 0$, and any cost distribution F,*

$$\frac{V(\boldsymbol{v})}{\mathsf{OFF}(\boldsymbol{v})} \geq \frac{n}{2n-1}. \tag{11}$$

Proof. Let $G_k \triangleq V(\boldsymbol{v}_{k:n})$ denote the greedy continuation value from the k-th worker onwards, with $G_{n+1} = 0$. By Lemma 1, $G_k = G_{k+1} + \psi(v_k - G_{k+1})$ for each $k \in [n]$, and $G_1 = V(\boldsymbol{v})$.

For any constant $c \geq 0$ and random variables $X_i = (v_i - Z_i)^+$, the pointwise bound $\max_i X_i \leq c + \sum_i (X_i - c)^+$ holds. Setting $c = G_2 \geq 0$ and taking expectations, the identity $\mathbb{E}[((v_i - Z_i)^+ - c)^+] = \psi(v_i - c)$ yields

$$\mathsf{OFF}(\boldsymbol{v}) \leq G_2 + \psi(v_1 - G_2) + \sum_{j=2}^{n} \psi(v_j - G_2).$$

The first two terms equal G_1 by the recursion for G_1. Defining $S_2 \triangleq \sum_{j=2}^{n} \psi(v_j - G_2)$,

$$\mathsf{OFF}(\boldsymbol{v}) \leq G_1 + S_2.$$

We derive two upper bounds on S_2. First, since $v_1 \geq v_j$ for all $j \geq 2$ and ψ is non-decreasing,

$$S_2 \leq (n-1)\,\psi(v_1 - G_2).$$

Second, define $S_k = \sum_{j=k}^{n} \psi(v_j - G_k)$ for $k \geq 2$, with $S_{n+1} = 0$. Since $G_k \geq G_{k+1}$, monotonicity of ψ gives $\psi(v_j - G_k) \leq \psi(v_j - G_{k+1})$ for $j \geq k+1$. Separating the leading term,

$$S_k \leq \psi(v_k - G_k) + S_{k+1} = (G_k - G_{k+1}) + S_{k+1}.$$

Telescoping from $k = 2$ to n yields $S_2 \leq G_2$.

Combining $(n-1)$ copies of $S_2 \le G_2$ with one copy of $S_2 \le (n-1)\,\psi(v_1 - G_2)$,

$$n \cdot S_2 \le (n-1)\big[G_2 + \psi(v_1 - G_2)\big] = (n-1)\,G_1,$$

so $S_2 \le \frac{n-1}{n}\,G_1$. Substituting back,

$$\mathsf{OFF}(\boldsymbol{v}) \le G_1 + \tfrac{n-1}{n}\,G_1 = \tfrac{2n-1}{n}\,G_1,$$

completing the proof. □

As $n \to \infty$, the lower bound $n/(2n-1)$ converges to $1/2$. To establish that this finite-horizon bound is tight, we evaluate greedy ranking and the optimal online policy $\mathsf{OPT}(\boldsymbol{v})$ under a Bernoulli distribution.

Let the cost $Z \in \{0,1\}$ follow a Bernoulli distribution, where $p = \mathbb{P}(Z=1)$ denotes the probability of a high cost realization. We construct a sequence of observable values:

$$v_k = 1 + (n-k)(1-p), \quad \text{for } k \in \{1,\dots,n\}.$$

Evaluating greedy ranking via backward induction on this sequence yields $V(\boldsymbol{v}) = n(1-p)$. The online optimum is achieved by deferring the highest-value worker to the final position, yielding:

$$\mathsf{OPT}(\boldsymbol{v}) = (n-1)(1-p) + n(1-p)p^{n-1}.$$

The approximation ratio for this instance is:

$$\frac{V(\boldsymbol{v})}{\mathsf{OPT}(\boldsymbol{v})} = \frac{n(1-p)}{(n-1)(1-p) + n(1-p)p^{n-1}} = \frac{n}{(n-1) + np^{n-1}}.$$

Taking the limit as $p \to 1$:

$$\lim_{p\to 1} \frac{V(\boldsymbol{v})}{\mathsf{OPT}(\boldsymbol{v})} = \frac{n}{(n-1)+n} = \frac{n}{2n-1}.$$

This establishes that the $n/(2n-1)$ bound is tight for every finite n. Since $\mathsf{OPT}(\boldsymbol{v}) \le \mathsf{OFF}(\boldsymbol{v})$, the same ratio simultaneously governs the performance of greedy ranking relative to the offline prophet.

6. Concluding Remarks and Future Work

This paper studies the sequential offering problem faced by on-demand platforms, where the sequence of workers and the dynamic wage trajectory must be jointly optimized. By identifying a location-shift structure arising from the observable heterogeneity among workers, we bridge the analytical gap between trivial i.i.d. environments and computationally intractable, arbitrarily heterogeneous settings.

Our main result establishes that when the unobservable cost distribution exhibits a non-increasing and convex density, the exact optimal policy is structurally simple: the platform should rank workers greedily by their observable net values and optimize wages dynamically via backward induction. In these settings, attempting to exploit dynamic pricing by sending initial low offers to inferior matches is strictly suboptimal.

Beyond these sufficient conditions for exact optimality, we demonstrate the robustness of this greedy heuristic through numerical experiments and additional theoretical results. We find strong empirical evidence for the optimality greedy ranking across several relevant distributions, even providing computer-aided proofs for specific cases. Moreover, even under suboptimal environments, the empirical optimality gap is surprisingly small, and we further establish tight worst-case guarantees for the performance of this greedy heuristic.

Several promising directions remain for future research. While our model focuses on sequential search for a single job, extending this framework to incorporate multiple jobs, parallel offerings, or correlated values would capture additional operational complexities. Furthermore, integrating online learning into the model to account for shifts in the platform's belief over the cost distribution F presents a valuable avenue for real-world platform operations (see Appendix C).

References

Agrawal S, Sethuraman J, Zhang X (2020) On optimal ordering in the optimal stopping problem. *Proceedings of the 21st ACM Conference on Economics and Computation*, 187–188.

Allon G, Cohen MC, Sinchaisri WP (2023) The impact of behavioral and economic drivers on gig economy workers. *Manufacturing & Service Operations Management* 25(4):1376–1393.

Amazon Flex (2023) Us — your guide to amazon flex offers. URL `https://flex.amazon.com/blog/2023/your-guide-to-amazon-flex-offers`, accessed on February 26, 2026.

Besbes O, Castro F, Lobel I (2021) Surge pricing and its spatial supply response. *Management Science* 67(3):1350–1367.

Beyhaghi H, Golrezaei N, Leme RP, Pál M, Sivan B (2021) Improved revenue bounds for posted-price and second-price mechanisms. *Operations research* 69(6):1805–1822.

Bimpikis K, Candogan O, Saban D (2019) Spatial pricing in ride-sharing networks. *Operations Research* 67(3):744–769.

Bloomberg (2021) Doordash workers are trying to game the algorithm to increase pay. URL `https://www.bloomberg.com/news/articles/2021-04-06/doordash-workers-are-trying-to-game-the-algorithm-to-increase-pay`, accessed on February 23, 2026.

Cachon GP, Daniels KM, Lobel R (2017) The role of surge pricing on a service platform with self-scheduling capacity. *Manufacturing & Service Operations Management* 19(3):368–384.

CareRev (2023) Smart rates: Dynamic pricing for hospital staffing. `https://www.carerev.com/solutions/smart-rates-healthcare-staffing`, accessed: 2026-02-26.

Castillo JC, Knoepfle D, Weyl G (2017) Surge pricing solves the wild goose chase. *Proceedings of the 2017 ACM Conference on Economics and Computation*, 241–242.

Correa J, Foncea P, Pizarro D, Verdugo V (2019) From pricing to prophets, and back! *Operations Research Letters* 47(1):25–29.

Derakhshan M, Ghasemi A, MacRury C (2026) Approximation algorithms for action-reward query-commit matching. *Working Paper* Available at `https://arxiv.org/abs/2603.13487`.

DoorDash (2025) How is your dasher paid? URL https://help.doordash.com/consumers/s/article/How-do-Dasher-earnings-work?language=en_US, accessed on February 18, 2026.

Epstein B, Ma W (2024) Selection and ordering policies for hiring pipelines via linear programming. *Operations Research* 72(5):2000–2013.

Garg N, Nazerzadeh H (2022) Driver surge pricing. *Management Science* 68(5):3219–3235.

Hill T (1983) Prophet inequalities and order selection in optimal stopping problems. *Proceedings of the American Mathematical Society* 88(1):131–137.

Krengel U, Sucheston L (1977) Semiamarts and finite values. *Bulletin of the American Mathematical Society* 83(4):745–747.

Ma H, Fang F, Parkes DC (2022) Spatio-temporal pricing for ridesharing platforms. *Operations Research* 70(2):1025–1041.

Peng B, Tang ZG (2022) Order selection prophet inequality: From threshold optimization to arrival time design. *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, 171–178 (IEEE).

Pollner T, Roghani M, Saberi A, Wajc D (2024) Improved online contention resolution for matchings and applications to the gig economy. *Math. Oper. Res.* .

Purohit M, Gollapudi S, Raghavan M (2019) Hiring under uncertainty. *International Conference on Machine Learning*, 5181–5189 (PMLR).

Reddit (2023) Really curious about people with low acceptance. URL https://www.reddit.com/r/UberEatsDrivers/comments/1s3qr3d/really_curious_about_people_with_low_acceptance/, accessed on May 4, 2026.

Samuel-Cahn E (1984) Comparison of threshold stop rules and maximum for independent nonnegative random variables. *the Annals of Probability* 1213–1216.

Taylor TA (2018) On-demand service platforms. *Manufacturing & Service Operations Management* 20(4):704–720.

Uber (2024) How much do uber drivers make? URL https://www.uber.com/us/en/drive/how-much-drivers-make/?city=new-york, accessed on February 26, 2026.

This page is intentionally blank. Proper e-companion title page, with INFORMS branding and exact metadata of the main paper, will be produced by the INFORMS office when the issue is being assembled.

Appendix A: Proof for Logistic Distribution

In this section, we specifically study case of Logistic distribution, with cumulative distribution function $F(x) = \frac{1}{1+e^{-x}}$, which is highly relevant in discrete choice theory as the foundational probability model for Multinomial Logit (MNL) frameworks.

PROPOSITION EC.1. *Suppose the private cost Z follows a standard Logistic distribution. Then the swap difference is identically zero, $D(x,y) \equiv 0$ for all $x, y \in (-\infty, \infty)$. Consequently, any ranking sequence yields the exact same expected welfare, and the greedy policy is globally optimal.*

Proof. The integrated cumulative distribution function evaluated from the lower bound of the support is given by:

$$\psi(x) = \int_{-\infty}^{x} F(t)dt = \int_{-\infty}^{x} \frac{e^t}{1+e^t} dt = \ln(1+e^x)$$

Recall from Section 3.3 that the swap difference is evaluated directly as $D(x,y) = \psi(y - \psi(x)) + \psi(x) - \psi(x - \psi(y)) - \psi(y)$. We evaluate the nested arguments to solve for $D(x,y)$. Substituting $\psi(y)$ into the downstream margin's shifted argument yields:

$$x - \psi(y) = x - \ln(1+e^y) = \ln\left(\frac{e^x}{1+e^y}\right)$$

Applying the integrated CDF ψ to this shifted argument yields:

$$\psi(x - \psi(y)) = \ln\left(1 + \frac{e^x}{1+e^y}\right) = \ln\left(\frac{1+e^y+e^x}{1+e^y}\right) = \ln(1+e^x+e^y) - \ln(1+e^y)$$

By exact symmetry, evaluating the upstream worker's shifted margin yields:

$$\psi(y - \psi(x)) = \ln(1+e^x+e^y) - \ln(1+e^x)$$

Substituting these evaluated terms back into the swap difference equation gives:

$$\begin{aligned} D(x,y) = & [\ln(1+e^x+e^y) - \ln(1+e^x)] + \ln(1+e^x) \\ & - [\ln(1+e^x+e^y) - \ln(1+e^y)] - \ln(1+e^y) \end{aligned}$$

All terms cancel perfectly, yielding $D(x,y) = 0$ for all x and y. Because any adjacent swap leaves the expected continuation value mathematically unchanged, sequence ranking is irrelevant. $\square$

This result reveals a connection to discrete choice theory. Specifically, the nested continuation values exactly mirror the standard Multinomial Logit (MNL) inclusive value (log-sum) formula. This sequence indifference is a direct manifestation of the Independence of Irrelevant Alternatives (IIA) property inherent to the Logistic distribution, dictating that the expected platform payoff remains perfectly invariant regardless of the order in which offers are made.

Appendix B: A Unified Quantile Framework for Sequential Offering

In this section, we extend the framework to establish the exact optimality of the greedy ranking across a variety of platform objectives, including profit maximization. We abstract away the wage space and formulate the platform's sequential dynamic program in the quantile (acceptance probability) space.

B.1. The Generalized Cost Model

Rather than directly choosing a wage w, the platform effectively chooses a target acceptance probability $q \in [0,1]$, which corresponds to offering a wage of $w = F^{-1}(q)$. If the platform approaches the worker $i \in I$ with a target acceptance probability q, the expected reward is:

$$R_i(q) = v_i q - K(q). \tag{EC.1}$$

The function $K(q)$ represents the expected cost incurred by the platform to induce an acceptance probability q. We consider three primary operational objectives:

1. Welfare Maximization: The platform seeks to maximize total generated surplus $v_i - Z_i$. The expected cost is the worker's expected private cost conditional on acceptance, $K_W(q) = \int_0^q F^{-1}(s)ds$.
2. Profit Maximization: The platform strictly maximizes its expected revenue minus the wage paid. The cost is the expected wage expenditure, $K_\Pi(q) = qF^{-1}(q)$.
3. Weighted Objective: Platforms frequently face regulatory or operational pressure to balance profit with worker welfare using a weight parameter $\alpha \in [0,1]$, $K_\alpha(q) = \alpha K_\Pi(q) + (1-\alpha)K_W(q)$.

To ensure the optimal dynamic pricing policies are mathematically well-defined, we formalize the requirements on the cost function:

Assumption EC.1. *The cost function $K(q)$ is strictly convex (i.e., $K''(q) > 0$) and four times continuously differentiable on $q \in (0,1)$.*

Remark: For welfare maximization, $K''_W(q) = 1/f(w) > 0$ holds naturally since $F(w)$ is strictly increasing. For profit maximization, strict convexity of $K_\Pi(q)$ is mathematically equivalent to assuming the virtual cost function $w + F(w)/f(w)$ is strictly increasing, which aligns perfectly with standard regularity assumptions in mechanism design.

B.2. Optimal Quantiles via Conjugate Duality

Fix a sequence of workers $u \in \Pi(v)$. Let $V(u_{t+1:n})$ denote the optimal expected continuation value derived from the downstream sequence from step $t+1$ to n. At step t, the platform selects q to maximize the immediate expected payoff plus the continuation value:

$$V(u_{t:n}) = V(u_{t+1:n}) + \max_{q\in[0,1]}\{q(u_t - V(u_{t+1:n})) - K(q)\} \tag{EC.2}$$

Notice that the optimal quantile depends entirely on the margin $x \triangleq u_t - V(u_{t+1:n})$. We define the generalized incremental surplus function $\psi(x)$ as the convex conjugate (Legendre-Fenchel dual) of the cost function $K(q)$:

$$\psi(x) \triangleq \max_{q\in[0,1]} \{qx - K(q)\} \tag{EC.3}$$

Remark: For the welfare objective $K_W(q)$, this dual formulation perfectly recovers the surplus function $\psi(x) = \int_{-\infty}^{x} F(s)ds$ utilized in Section 3.

By Assumption EC.1, the strict convexity of $K(q)$ guarantees a unique optimal probability $q^*(x)$ characterized by the first-order condition $x = K'(q)$ (restricting to margins where the unconstrained probability is interior). By standard conjugate duality (and the Envelope Theorem), the derivative of the surplus function evaluates exactly to this optimal acceptance probability:

$$\psi'(x) = q^*(x) = (K')^{-1}(x) \tag{EC.4}$$

Consequently, the dynamic programming recursion cleanly simplifies to $V(u_{t:n}) = V(u_{t+1:n}) + \psi(u_t - V(u_{t+1:n}))$, perfectly mirroring the structure established in Lemma 1.

B.3. Global Optimality of Greedy Ranking

To prove that the greedy ranking (sorting by descending order of u_t) is globally optimal, we consider the expected welfare difference between an arbitrary sequence and a sequence where an adjacent inverted pair is swapped. Let x and y be the margins of the downstream and upstream workers, respectively ($x \le y$). As established in Equation (7), the expected difference is:

$$D(x,y) \triangleq \psi(y - \psi(x)) + \psi(x) - \psi(x - \psi(y)) - \psi(y) \tag{EC.5}$$

As proven in Theorem 1, global optimality ($D(x,y) \ge 0$) is guaranteed if the expected acceptance probability function $\psi'(x)$ is concave ($\psi'''(x) \le 0$) and its derivative is convex ($\psi^{(4)}(x) \ge 0$). By evaluating the derivatives of the inverse function $\psi'(x) = (K')^{-1}(x)$, we map this requirement directly to the structural properties of $K(q)$.

THEOREM EC.1. *For any objective parameterized by a strictly convex cost function $K(q)$, the greedy ranking is globally optimal if $K(q)$ satisfies the following conditions for all $q \in (0,1)$:*

1. $K'''(q) \ge 0$
2. $3(K'''(q))^2 - K^{(4)}(q)K''(q) \ge 0$

Proof. Applying the Inverse Function Theorem to $\psi'(x) = (K')^{-1}(x)$ yields $\psi''(x) = 1/K''(q)$. Applying the chain rule $\frac{dq}{dx} = \psi''(x)$ to subsequent derivatives yields:

$$\psi'''(x) = \frac{-K'''(q)}{(K''(q))^3}$$
$$\psi^{(4)}(x) = \frac{3(K'''(q))^2 - K^{(4)}(q)K''(q)}{(K''(q))^5}$$

Since strict convexity dictates $K''(q) > 0$, the denominators are strictly positive. Thus, $\psi'''(x) \leq 0$ holds if and only if $K'''(q) \geq 0$, and $\psi^{(4)}(x) \geq 0$ holds if and only if the numerator $3(K'''(q))^2 - K^{(4)}(q)K''(q) \geq 0$. $\square$

The quantile formulation provides a universal test for global optimality, allowing us to map the conditions of Theorem EC.1 directly back to the physical probability density function $f(w)$. Note that derivatives with respect to q can be mapped to derivatives with respect to w via the chain rule $\frac{dw}{dq} = \frac{1}{f(w)}$.

Profit Maximization. For the profit cost function $K_\Pi(q) = qF^{-1}(q)$, the first derivative recovers the standard virtual cost formulation $K'_\Pi(q) = w + F(w)/f(w)$. The higher-order derivatives yield:

$$K''_\Pi(q) = \frac{2f(w)^2 - F(w)f'(w)}{f(w)^3}$$
$$K'''_\Pi(q) = \frac{-F(w)f(w)f''(w) + 3F(w)(f'(w))^2 - 3f(w)^2 f'(w)}{f(w)^5}$$

Proposition EC.2. *For a profit-maximizing platform, if the cost distribution exhibits a non-increasing ($f'(w) \leq 0$) and log-concave ($f(w)f''(w) \leq (f'(w))^2$) density, then the effective acceptance probability is strictly concave (Condition 1 holds).*

Proof. Log-concavity bounds the second derivative such that $-F(w)f(w)f''(w) \geq -F(w)(f'(w))^2$. Since $F(w) \geq 0$ and $f'(w) \leq 0$, the remaining terms in the numerator of $K'''_\Pi(q)$, specifically $2F(w)(f'(w))^2 - 3f(w)^2 f'(w)$, are unconditionally non-negative. Summing these bounds ensures $K'''_\Pi(q) \geq 0$. $\square$

Remark: It is intuitively tempting to conjecture that log-concavity also perfectly guarantees the fourth-order global sufficiency condition (Condition 2). However, evaluating $K^{(4)}_\Pi(q)$ introduces the third derivative of the density, $f'''(w)$, which is unbounded by standard log-concavity. Thus, global optimality must be verified exactly for canonical distributions.

Theorem EC.2. *For a profit-maximizing platform, if the unobservable private costs follow a Uniform or Exponential distribution, the greedy ranking is exactly globally optimal.*

Proof. For the Uniform distribution ($f(w) = 1/a, F(w) = w/a$), the virtual cost is $K'_\Pi(q) = 2w$. Subsequent derivatives yield $K''_\Pi(q) = 2a$, and identically $K'''_\Pi(q) = K^{(4)}_\Pi(q) = 0$. This trivially satisfies both sufficiency conditions of Theorem EC.1.

For the Exponential distribution ($f(w) = \lambda e^{-\lambda w}$), the derivatives with respect to q evaluate precisely to:

$$K''_\Pi(q) = \lambda^{-1}\left(e^{\lambda w} + e^{2\lambda w}\right)$$
$$K'''_\Pi(q) = \lambda^{-1}\left(e^{2\lambda w} + 2e^{3\lambda w}\right)$$
$$K^{(4)}_\Pi(q) = \lambda^{-1}\left(2e^{3\lambda w} + 6e^{4\lambda w}\right)$$

Condition 1 ($K_\Pi''' \geq 0$) holds trivially. Evaluating the expression for Condition 2 yields:

$$\begin{aligned} 3(K_\Pi'''(q))^2 - K_\Pi^{(4)}(q)K_\Pi''(q) &= \lambda^{-2}\Big[3\left(e^{2\lambda w}+2e^{3\lambda w}\right)^2 - \left(2e^{3\lambda w}+6e^{4\lambda w}\right)\left(e^{\lambda w}+e^{2\lambda w}\right)\Big] \\ &= \lambda^{-2}\Big[\left(3e^{4\lambda w}+12e^{5\lambda w}+12e^{6\lambda w}\right) - \left(2e^{4\lambda w}+8e^{5\lambda w}+6e^{6\lambda w}\right)\Big] \\ &= \lambda^{-2}\Big[e^{4\lambda w}+4e^{5\lambda w}+6e^{6\lambda w}\Big] \geq 0 \end{aligned}$$

Because both conditions are strictly satisfied everywhere, the Exponential distribution geometrically guarantees exact global optimality. □

Appendix C: Robustness to Deterministically Decaying Markets

In our baseline model, we assumed that the reservation cost distribution F remains stationary across all positions in the sequence. In practice, however, on-demand platforms often deploy predefined operational pipelines to handle sequential rejections. For instance, when an initial offer is rejected, the platform may systematically relax geographic matching radii or extend the search to lower-tier worker pools. Even without explicit platform interventions, the sheer passage of time after multiple rejections naturally shrinks the remaining delivery window, effectively increasing the remaining workers' private fulfillment costs.

To clearly model the temporal progression of such a pipeline, let $t \in [n]$ denote the forward index of the sequence of offers. At step t, the approached worker draws a private cost from a time-specific distribution F_t. We capture the systematic escalation of expected costs by modeling a *deterministically decaying market*, formalized via First-Order Stochastic Dominance (FOSD) across the sequence of steps.

Definition EC.1 (Deterministically Decaying Market). A sequential offering market is deterministically decaying if the unobservable private costs stochastically increase over time. Equivalently, the acceptance probabilities systematically decay, satisfying First-Order Stochastic Dominance (FOSD): $F_1(w) \geq F_2(w) \geq \cdots \geq F_n(w)$ for all $w \in \mathbb{R}$.

Intuitively, deteriorating acceptance probabilities might seem to break the greedy heuristic, potentially incentivizing the platform to strategically "save" the highest-value workers for later stages to offset the worsening market conditions. Surprisingly, our next result establishes that the greedy ranking remains universally optimal. Furthermore, the mathematical mechanics reveal that the systematic decay completely absorbs the structural shape constraints entirely in the initial step of the pipeline.

Theorem EC.3. *Consider a deterministically decaying market over $n \geq 2$ steps with a worker pool of observable net values $v_1 \geq v_2 \geq \cdots \geq v_n > 0$. Assume the initial cost distribution F_1 has a support bounded below. If for every downstream step $t \in \{2, 3, \ldots, n\}$, the cost distribution F_t satisfies the* *baseline stability conditions (i.e., its CDF is concave on its support and admits a convex probability density function), then the greedy sequence $\boldsymbol{v} \in \Pi(\boldsymbol{v})$ is globally optimal. Crucially, the initial step's distribution F_1 requires no shape constraints whatsoever.*

Proof. We prove global optimality using an adjacent-swap argument evaluated via backward induction. Following Lemma 1, let $V(\boldsymbol{u}_{t:n})$ denote the optimal expected continuation value from step t onward, with the terminal condition $V(\boldsymbol{u}_{n+1:n}) = 0$. For any sequence $u \in \Pi(\boldsymbol{v})$, the dynamic programming update at step t is:

$$V(\boldsymbol{u}_{t:n}) = V(\boldsymbol{u}_{t+1:n}) + \psi_t(u_t - V(\boldsymbol{u}_{t+1:n})) \tag{EC.6}$$

where $\psi_t(x) = \int_{-\infty}^{x} F_t(s)ds$. Since F_1 has a support bounded below by some constant L, $F_1(w) = 0$ for all $w \leq L$. By the FOSD condition, $0 = F_1(w) \geq F_t(w) \geq 0$ inherently guarantees $F_t(w) = 0$ for all $w \leq L$ and all $t \in [n]$. Thus, all distributions natively share a uniformly bounded lower support, ensuring ψ_t is well-defined.

Consider an arbitrary non-greedy sequence $u \in \Pi(\boldsymbol{v}) \setminus \{v\}$. There must exist at least one adjacent inverted pair of workers at steps t and $t+1$ (where $t \in [n-1]$) whose observable net values satisfy $u_t < u_{t+1}$.

Let $z = V(\boldsymbol{u}_{t+2:n})$ denote the optimally extracted continuation value from the remaining sequence (steps $t+2$ through n). Importantly, z depends exclusively on the set of remaining downstream workers and their respective distributions, making it perfectly invariant to the local ordering of the workers at steps t and $t+1$. We define the margins of our inverted pair over this downstream continuation value as $x = u_t - z$ and $y = u_{t+1} - z$. Because the pair is inverted, it strictly holds that $x < y$.

If the platform evaluates the pair in their current non-greedy order, the expected continuation value entering step t is:

$$V_{orig} = z + \psi_{t+1}(y) + \psi_t(x - \psi_{t+1}(y)). \tag{EC.7}$$

If we execute a local swap to place the superior worker first (greedy order), the re-optimized value becomes:

$$V_{swap} = z + \psi_{t+1}(x) + \psi_t(y - \psi_{t+1}(x)). \tag{EC.8}$$

We isolate the local welfare impact by defining the deterministic swap difference $D_t(x, y) = V_{swap} - V_{orig}$. Canceling the constant z yields:

$$D_t(x, y) = [\psi_t(y - \psi_{t+1}(x)) + \psi_{t+1}(x)] - [\psi_t(x - \psi_{t+1}(y)) + \psi_{t+1}(y)]. \tag{EC.9}$$

To prove that swapping the workers weakly improves the continuation value, we must establish $D_t(x, y) \geq 0$ for all $x \leq y$. Because F_t and F_{t+1} are cumulative distribution functions bounded in

$[0,1]$, their integrals ψ_t and ψ_{t+1} are 1-Lipschitz. The composition of Lipschitz functions is Lipschitz, meaning $D_t(x,y)$ is absolutely continuous and differentiable almost everywhere. We evaluate its partial derivative with respect to x:

$$\frac{\partial D_t(x,y)}{\partial x} = F_{t+1}(x)\left[1 - F_t(y - \psi_{t+1}(x))\right] - F_t(x - \psi_{t+1}(y)). \tag{EC.10}$$

Let $D^*_{t+1}(x,y)$ define the purely stationary swap difference if *both* steps t and $t+1$ operated strictly under the downstream distribution F_{t+1}. By exact symmetry, its partial derivative is:

$$\frac{\partial D^*_{t+1}(x,y)}{\partial x} = F_{t+1}(x)\left[1 - F_{t+1}(y - \psi_{t+1}(x))\right] - F_{t+1}(x - \psi_{t+1}(y)). \tag{EC.11}$$

By the deterministically decaying market condition, FOSD guarantees $F_t(w) \geq F_{t+1}(w)$ for all $w \in \mathbb{R}$. First, because $F_{t+1}(x)$ is a probability, $F_{t+1}(x) \geq 0$. Multiplying the FOSD inequality evaluated at the leading argument by $-F_{t+1}(x)$ bounds the first term:

$$-F_{t+1}(x)F_t(y - \psi_{t+1}(x)) \leq -F_{t+1}(x)F_{t+1}(y - \psi_{t+1}(x)). \tag{EC.12}$$

Second, evaluating the trailing argument directly dictates $-F_t(x - \psi_{t+1}(y)) \leq -F_{t+1}(x - \psi_{t+1}(y))$. Summing these bounded components establishes a strict global inequality on the partial derivatives almost everywhere:

$$\frac{\partial D_t(x,y)}{\partial x} \leq \frac{\partial D^*_{t+1}(x,y)}{\partial x} \quad \text{for all } x \leq y. \tag{EC.13}$$

At the boundary where margins are identical ($x = y$), the swap differences trivially equal zero: $D_t(y,y) = 0$ and $D^*_{t+1}(y,y) = 0$. By the Fundamental Theorem of Calculus for Lebesgue integrals, integrating the negative partial derivatives backward from y down to x yields:

$$D_t(x,y) = \int_x^y \left(-\frac{\partial D_t(s,y)}{\partial s}\right) ds \geq \int_x^y \left(-\frac{\partial D^*_{t+1}(s,y)}{\partial s}\right) ds = D^*_{t+1}(x,y). \tag{EC.14}$$

Since $t \in [n-1]$, the downstream step $t+1 \geq 2$. Thus, F_{t+1} satisfies the baseline stability conditions (concave CDF, convex density) by assumption, and Theorem 1 guarantees the stationary baseline is stable: $D^*_{t+1}(x,y) \geq 0$. By transitivity, $D_t(x,y) \geq 0$. Executing the local adjacent swap weakly increases the expected continuation value entering step t.

Finally, we must prove that this local improvement monotonically propagates backward to increase the total sequence welfare. For any preceding step $k \leq t$, we analyze the sensitivity of the step k continuation value with respect to its downstream continuation value $V \triangleq V(\boldsymbol{u}_{k+1:n})$. Because the derivative of the surplus function is exactly the acceptance probability $\psi'_k(w) = F_k(w) \in [0,1]$ almost everywhere, ψ_k is Lipschitz continuous with constant 1. Therefore, for any finite increase $\Delta > 0$ in the downstream continuation value:

$$[V + \Delta + \psi_k(u_k - (V + \Delta))] - [V + \psi_k(u_k - V)] = \Delta - [\psi_k(u_k - V) - \psi_k(u_k - V - \Delta)] \geq \Delta - \Delta = 0. \tag{EC.15}$$

This mathematically guarantees that the dynamic programming recursion is monotonically non-decreasing. The local welfare gain at step t weakly increases $V(\boldsymbol{u}_{t-1:n})$, chaining completely up the sequence to weakly increase the total expected welfare $V(\boldsymbol{u}_{1:n})$. By iteratively identifying and swapping adjacent inverted pairs in this manner, any arbitrary ranking sequence can be transformed into the greedy sequence without ever decreasing expected welfare. This establishes the global optimality of the greedy ranking. □

Appendix D: Methodology for Numerical Certification

To bridge the gap between empirical evidence from Section 4 and absolute mathematical certainty, we formalize a rigorous *computer-assisted* proof.

This framework transforms an uncountable, unbounded, and degenerate continuous space into a finite, bounded grid where global optimality can be computationally certified. The methodology proceeds in three phases:

1. We truncate the infinite 2D space into a bounded region.
2. We establish certified "safe zones" around the origin and diagonal, clearing the boundaries where standard floating-point solvers fail.
3. We derive bounds for the partial derivatives of $D(x,y)$ to formally guarantee global continuous optimality from finite discrete evaluations.

D.1. Proofs.

D.1.1. Domain Compactification

Proposition EC.3. *For any continuous cost distribution F with non-negative support ($F(z)=0$ for $z\le 0$), the swap difference satisfies $D(x,y)\ge 0$ unconditionally for all $x\le\psi(y)$.*

Proof. If $x\le\psi(y)$, the lagged argument $x-\psi(y)$ is non-positive. Because private costs are non-negative, $F(z)=0$ and $\psi(z)=0$ for all $z\le 0$. The swap difference reduces to $D(x,y)=\psi(y-\psi(x))+\psi(x)-\psi(y)$. By the Fundamental Theorem of Calculus, $\psi(y)-\psi(y-\psi(x))=\int_{y-\psi(x)}^{y}F(t)dt$. Because the CDF is strictly bounded by $F(t)\le 1$, this integral is strictly bounded above by the integration interval width, $\psi(x)$. Rearranging yields $\psi(x)-(\psi(y)-\psi(y-\psi(x)))\ge 0$, establishing $D(x,y)\ge 0$. □

Proposition EC.4. *For distributions with unbounded left support spanning $(-\infty,\infty)$ (e.g., Normal, Logistic), there exists an exact boundary curve $X_{up}(y)$ such that $D(x,y)>0$ for all $x\le X_{up}(y)$.*

Proof. Rewriting the swap difference exactly as the difference of two integrals yields $D(x,y)=\int_{x-\psi(y)}^{x}F(t)dt-\int_{y-\psi(x)}^{y}F(t)dt$. Bounding the subtrahend integral strictly by its maximum possible rectangular area $\psi(x)F(y)$ provides a strict lower envelope:

$$D(x,y)>\psi(x)\left[1-F(y)-\frac{\psi(x-\psi(y))}{\psi(x)}\right] \tag{EC.16}$$

For distributions with log-concave left tails, applying L'Hôpital's Rule yields $\lim_{x\to-\infty}\frac{\psi(x-c)}{\psi(x)} = \lim_{x\to-\infty}\frac{F(x-c)}{F(x)} = \lim_{x\to-\infty}\frac{f(x-c)}{f(x)} = 0$. Because $1-F(y)>0$ for any finite y, the bracketed term is mathematically guaranteed to become strictly positive in the deep tail. The root where this envelope crosses zero defines the exact analytical boundary $X_{up}(y)$, strictly compactifying domains unbounded below. $\square$

Proposition EC.5. *Let $\mu = \mathbb{E}[Z] < \infty$. As $y\to\infty$, the swap difference converges uniformly to a 1D limit function: $L(c) = \psi(\mu+c) - c - \psi(\mu-c)$ where $c = y-x$. Uniform convergence guarantees a finite upper truncation limit Y_{max} if and only if $\mu > m$ (Mean strictly greater than Median).*

Proof. Evaluating the derivative of the limit function as the gap $c\to 0$ yields $L'(0) = F(\mu+0) - 1 + F(\mu-0) = 2F(\mu)-1$. Because $L(0)=0$, the tail limit remains strictly positive only if $L'(0)>0$, structurally requiring $F(\mu)>1/2$. By definition of the median m, this dictates $\mu>m$. If this topological condition is met, there exists a uniform finite threshold Y_{max} beyond which $D\geq 0$ globally. $\square$

D.1.2. Boundary "safe zones"

Proposition EC.6. *Let $S(y) \triangleq F(y-\psi(y))[1+F(y)] - F(y)$. Let f_{sup} be the maximum density locally bounded within the interval $z\in[y-\epsilon, y]$. If $S(y)>0$, then $D(x,y)>0$ for all $x\in(y-\epsilon(y), y)$, where $\epsilon(y) = \frac{2S(y)}{f_{sup}}$.*

Proof. Expanding $D(y-\Delta, y)$ around the gap $\Delta = 0$ using a Taylor series with a Lagrange remainder yields $D \geq \Delta S(y) + \frac{\Delta^2}{2}\inf\frac{\partial^2 D}{\partial x^2}$. The second derivative is bounded strictly below by $-f(x-\psi(y)) \geq -f_{sup}$. Setting the lower bound $\Delta S(y) - \frac{\Delta^2}{2} f_{sup} > 0$ analytically clears the diagonal interval. Iterative fixed-point derivation of f_{sup} allows clearance even for distributions with global singularities ($f_{max} = \infty$). $\square$

D.1.3. Lipschitz Bounding

Lemma EC.1. *For any two-dimensional cell defined by $[x_{min}, x_{max}] \times [y_{min}, y_{max}]$ evaluated over the sequence domain $x\leq y$, the absolute gradient magnitudes across the entire continuous cell volume are bounded by:*

$$\left|\frac{\partial D}{\partial x}\right| \leq F(x_{max}) \tag{EC.17}$$

$$\left|\frac{\partial D}{\partial y}\right| \leq \max\Big(\min\big[F(y_{max}),\, \psi(x_{max}) f_{max}\big],\, F(x_{max})\Big) \tag{EC.18}$$

Proof. Evaluating the exact continuous X-derivative yields $\frac{\partial D}{\partial x} = F(x)[1-F(y-\psi(x))] - F(x-\psi(y))$. The leading positive component is bounded above by $F(x)\leq F(x_{max})$, and the subtracted component ensures the lower bound is trivially regulated $\geq -F(x_{max})$, yielding the bound $|\partial D/\partial x| \leq F(x_{max})$.

Evaluating the exact Y-derivative yields $\frac{\partial D}{\partial y} = -[F(y) - F(y - \psi(x))] + [F(y)F(x - \psi(y))]$. Define these non-negative components as A and B. By absolute value algebra, $|-A + B| \le \max(A, B)$. B is structurally bounded by $1 \cdot F(x - \psi(y)) \le F(x_{max})$. A is trivially bounded by $F(y) \le F(y_{max})$. Furthermore, by the exact Mean Value Theorem, the CDF difference A over the interval $\psi(x)$ is strictly bounded by $\psi(x)f(\xi) \le \psi(x_{max})f_{max}$. Taking the most restrictive valid bound yields $\sup A \le \min(F(y_{max}), \psi(x_{max})f_{max})$. Taking the maximum between the bounds for A and B completes the proof. □

THEOREM EC.4. *By Propositions EC.3–EC.6, we define a strictly interior, compact set $\mathcal{K}$. Partition $\mathcal{K}$ using an Adaptive Quadtree grid. By the Multivariate Mean Value Theorem and Lemma EC.1, the true continuous function $D(x, y)$ is mathematically guaranteed to be strictly positive everywhere inside a cell if its center evaluation $D(x_c, y_c)$ satisfies:*

$$D(x_c, y_c) > \left(\frac{x_{max} - x_{min}}{2}\right) \left|\frac{\partial D}{\partial x}\right|_{sup} + \left(\frac{y_{max} - y_{min}}{2}\right) \left|\frac{\partial D}{\partial y}\right|_{sup} \tag{EC.19}$$

Cells failing this strict mathematical hurdle are recursively subdivided.

D.2. Certifiable Distributions

Because this unified CAP framework relies on cumulative distributions, mean-value bounds, and real analysis rather than heuristics, it natively absorbs unbounded domains, zero-origin topological decay, and infinite singularities. It mathematically certifies the following:

1. Half-Normal Distribution: Features $f(0) \approx 0.798 > 0$ and $\mu \approx 0.798 > m \approx 0.674$. Certified definitively up to its analytical tail truncation $Y_{max} = 10.0$.
2. Log-normal Distribution: Evaluated under $\mu_{base} = 0.0, \sigma = 0.5$. It features a flat origin $f(0) = 0$, triggering catastrophic numerical $\mathcal{O}(x^3)$ decay. The framework natively averts infinite recursion via the Second-Order MVT bound in Lemma EC.1. Because the $\min(\cdot)$ operator dynamically activates $\psi(x_{max})f_{max}$, the algorithm mathematically shrinks the Y-radius directly to the true geometry without heuristic approximations.