Saving Lives in Animal Shelters: Strategic Time Allocation Under Capacity Constraints

Qiuxia (Katalia) Chen
Naveen Jindal School of Management, the University of Texas at Dallas, Qiuxia.Chen@utdallas.edu

Ernan Haruvy
Desautels Faculty of Management, McGill University ernan.haruvy@mcgill.ca

Elena Katok
Naveen Jindal School of Management, the University of Texas at Dallas, ekatok@utdallas.edu

Each year, approximately 425,000 animals are euthanized in U.S. shelters, many due to operational constraints rather than health issues. Municipal shelters face unique challenges: legal mandates require accepting all animals regardless of capacity, while existing operations research models assume homogeneous populations or controllable demand—assumptions that fail to capture shelter realities where dogs vary dramatically in adoption prospects. This study addresses how shelters can strategically allocate limited time across heterogeneous dog populations to maximize lifesaving outcomes under capacity constraints.

We develop an integrated framework combining queuing theory, survival analysis, and constrained optimization to optimize grace period allocation across dog types. Using 12.25 years of data from Dallas Animal Services—comprising over 200,000 dogs—we implement a four-stage approach: (1) queuing models to determine system-wide time constraints, (2) clustering analysis to segment dogs into operationally relevant types, (3) survival modeling to estimate time-dependent adoption probabilities, and (4) nonlinear optimization to reallocate time for maximum live release rates.

Counterfactual analysis demonstrates potential annual savings of 1,400 dogs through strategic reallocation, with validation showing live release rates increased from 78% to 87% during partial implementation starting in 2024. The framework provides implementable decision rules for grace period assignment based on observable dog characteristics. By treating time as a strategic resource rather than uniform constraint, managers achieve more humane outcomes without capacity expansion. The approach extends to other capacity-constrained public services facing heterogeneous populations, offering a generalizable framework for balancing operational efficiency with equity considerations in resource-scarce environments.

Keywords: Animal shelter operations, Queuing models, Survival analysis, Public sector service systems

1. Introduction

In 2024, over 5.8 million animals enter the U.S. shelter and rescue system through approximately 13,500 animal welfare organizations, including about 4,000 government-supported shelters and 9,500 rescues without a government contract (Shelter Animals Count, 2024). To manage this volume, shelters must routinely make life-and-death decisions under severe operational constraints. When intake exceeds capacity, many shelters resort to euthanizing healthy, adoptable animals—not due to behavioral or medical issues, but simply because physical space and time have run out. Estimates for the number of animals who died in U.S. shelters in 2024 vary between 425,000 (Best Friends Animal Society, 2024) and 750,000 (Shelter Animals Count, 2024), many euthanized due to operational constraints rather than health or behavioral issues.

Municipal shelters face a unique operational dilemma. Unlike private shelters that can limit intake when approaching capacity, municipal facilities operate under legal mandates to accept all animals regardless of available space. This creates a non-deferrable queue where excess demand cannot be rejected or postponed. The Guidelines for Standards of Care in Animal Shelters (The Association of Shelter Veterinarians, 2022) emphasize that capacity management requires careful balance between throughput and animal welfare. When intake exceeds capacity, managers must decide which dogs receive extended time to find homes, and which are euthanized to make room for new arrivals. These decisions typically rely on staff experience and informal rules, leading to potentially suboptimal outcomes. In this paper we propose a systematic method for improving these outcomes.

The consequences of potentially sub-optimal decisions are severe, because the average proportion of dogs that enter the US shelter system and are successfully saved declined from 88% in 2020 to 82% in 2024, representing approximately 425,000 animals killed in U.S. shelters in 2024 compared to 302,000 in 2020 (Best Friends Animal Society, 2024). Municipal shelters bear the greatest burden, with dog intake rising 3.1% at municipal facilities from 2023 to 2024 while transfers to private organizations decreased by nearly 30% (Castle, 2024).

The research question we address in this paper is how should shelters allocate finite time and space to maximize lifesaving outcomes across a heterogeneous dog population? We focus specifically on optimizing grace period—the length of time a dog remains in the shelter before facing euthanasia. We focus on the live release rate (*LLR*)—the share of animals adopted, transferred, or returned to owner—as the key performance metric, an industry-standard measure of shelter effectiveness (Weiss et al., 2013). While prior research has explored capacity management in animal shelters, existing models fail to capture two important realities of municipal shelter operations: the heterogeneity of adoption patterns across different types of dogs, and the inability to reject arrivals even when at capacity.

Several analytical frameworks have addressed animal shelter optimization challenges. Kang and Han (2019) developed a mathematical model to minimize euthanasia through improved inter-shelter transfers and space allocation, demonstrating a 30% reduction in euthanasia. Their approach, however, assumes coordination across multiple facilities—a luxury unavailable to isolated municipal shelters. Turken et al. (2021) applied queuing theory to animal shelter operations by treating the shelter as a system with limited capacity in which animals are turned away once all spaces are filled. This approach is appropriate for private rescues, but is unsuitable for municipal shelters, that typically cannot reject arrivals. Also, their model assumes homogeneous service times, while substantial heterogeneity in adoption probabilities across dog types exists. For example, small, healthy puppies often find homes within days, while large or medically complex dogs may remain in the system for weeks or months.

Other optimization efforts have focused on facility design and capacity estimation. Jalayer et al. (2024) proposed multi-criteria optimization for shelter layouts that reduce visual stressors. The University of Wisconsin–Madison Shelter Medicine Program (2023) provides calculators to estimate appropriate shelter capacity based on intake rates and staffing levels. While valuable, these approaches treat capacity as the primary lever rather than optimizing allocation within fixed constraints.

We develop an integrated framework that combines queuing theory, survival analysis, and constrained optimization to maximize the proportion of animals who leave the shelter alive. This approach follows earlier inventory models that explicitly incorporate replenishment and disposal policies under stochastic demand and capacity limits (Pinçe et al., 2008).

Using 12.25 years of operational data from Dallas Animal Services—one of the largest municipal shelters in the United States—we demonstrate how strategic time allocation can save thousands of additional dogs annually without expanding physical capacity. This framework treats time as a scarce resource to be optimized rather than a uniform constraint, recognizing that different types of dogs benefit differently from extended shelter stays.

Our work builds on recent advances in equity-aware resource allocation (Li et al., 2021; Gaur and Honhon, 2006) and extends queuing models for animal shelters beyond existing frameworks. We overcome the limitations of prior models by introducing a queuing framework that incorporates heterogeneous adoption patterns and enforces mandatory intake constraints. Although our primary focus is on municipal animal shelters, the underlying methodology generalizes to other sustainability and public service contexts—such as food waste reduction (Chehrazi et.al., 2025), food rescue (Benade and Alptekinoglu, 2024) and medical supply redistribution—where perishable, heterogeneous resources must be reallocated under strict operational constraints.

Our analysis reveals a counterintuitive insight: allocating more time to slower-adopting dogs (such as larger dogs) while reducing time for faster-adopting dogs (such as smaller breeds) can increase overall adoptions. This occurs because the marginal benefit of additional days varies dramatically across dog types—a nuance missed by uniform policies. Counterfactual analysis estimates potential annual savings of over 1000 additional dogs through strategic reallocation, with validation showing that live release rates increased from 87% to 92% during partial implementation of our approach starting in 2024.

This work contributes to both operations management theory and animal welfare practice. Theoretically, we extend queuing models to incorporate heterogeneous service populations under mandatory intake constraints. Practically, we provide implementable decision rules that translate complex analytics into daily operational guidance. While focused on animal shelters, our framework applies broadly to capacity-constrained public services facing heterogeneous demand—from homeless shelters to emergency medical care.

2. Data

We construct a panel dataset using operational records from Dallas Animal Services (DAS), that manages over 20,000 intakes annually, making it one of the largest municipal shelters in the United States. The dataset combines three primary data sources: (1) intake and outcome records containing dog-level characteristics (breed, size, health condition, intake type, outcome type, and timestamps) (2) daily shelter report cards documenting system-wide inventory counts, admissions, and exits; and (3) internal administrative notes used to resolve inconsistencies across reporting periods and classification changes.

Our sample spans all impounded dogs between January 1, 2013, and March 31, 2025, yielding 210,844 unique dogs with 259,228 individual visit records. The dataset captures the complete shelter experience for each dog, including multiple visits.

Because DAS did not consistently record daily shelter load, we reconstructed daily inventory levels (I_t) for 2013–2025 using the identity:

$$I_t = I_{t-1} + Intakes_t - Outcomes_t \tag{1}$$

We used known daily inventory values as anchor points: January 1 values for 2017–2019, October 3, 2019 for 2020–2024, and both January 1 and March 27, 2025 values for 2025. For 2013–2016, we back-propagated from the January 1, 2017 value because inventory counts for those years did not exist. This reconstruction assumes no exogenous shocks to data reporting and we validated it against known shelter census records.

For each intake *i,* the outcome is a binary variable called *positive outcome*, π_i defined as:

$$\pi_i = \begin{cases} 1 & \text{if adopted, transferred, or returned to owner} \\ 0 & \text{euthanized or died in care.} \end{cases}$$

We treat foster placement as a transitional rather than a final outcome. Dogs fostered and later adopted or euthanized are coded based on their final disposition. The system-level metric is the Live Release Rate (LLR), defined as:

$$LLR = \frac{\sum_{i \in N} \pi_i}{N} \tag{2}$$

where N is the total number of dogs the shelter managed during a specific time period. Dogs who move to foster homes are not included in the LLR calculation until their final outcome (usually adoption).

For dogs with multiple visits, we handle the data in two ways depending on the analysis. For survival analysis and clustering, we consolidate the data to one observation per dog using attributes from their most recent visit, while also tracking the total number of visits and using shelter time from the final stay; the details of this consolidation procedure are presented in on-line appendix A[1]. This approach allows us to model individual dog outcomes while accounting for each animal's full history in the shelter system. For system-level metrics and queuing analysis, we retain all visits to accurately capture shelter load and flow dynamics.

Table 1 presents summary statistics for our analytic sample of 259,228 intakes, showing substantial heterogeneity across size (40% large dogs), breed (23% bully breeds), age (25% under 6 months), and health status (79% healthy at intake). The average shelter stay is 7.47 days, with 75% achieving positive outcomes.

[1] Link for on-line appendix:
https://www.dropbox.com/scl/fi/q2ic6vljl4mqas4b7v9hp/Appendices_09162025KC.pdf?rlkey=97n6mnthdbq6eh1bu3ynv6x0q&st=i0rr8s46&dl=0

Table 1. Summary Statistics (N=245,728 intake records with final outcomes)

Variable Name	*Explanation*	*Data Source*	*Mean (SD)*
Number of visits	# time a dog visited the shelter	DAS outcome data	1.54 (1.05)
Toy	Size variables	DAS Outcome data	0.03 (0.18)
Small			0.22 (0.42)
Medium			0.33 (0.47)
Large			0.40 (0.49)
X-large			0.02 (0.13)
Female	Gender	DAS Outcome data	0.47 (0.50)
Bully	Primary breed includes	DAS Outcome data	0.23 (0.42)
Husky			0.02 (0.15)
Lab			0.11 (0.31)
Shepherd			0.11 (0.32)
Chihuahua			0.10 (0.30)
Black	Appearance/Color	DAS intake & outcome data	0.28 (0.45)
White		DAS intake & outcome data	0.17 (0.38)
Fluffy		Constructed with DAS outcome data	0.04 (0.20)
Age	Age of a dog at intake	DAS Outcome data	2.30 (2.49)
Baby	Younger under 6 months	DAS Outcome data	0.25 (0.43)
Surrendered	Owner surrendered	DAS Outcome data	0.27 (0.44)
Injured	Health condition	Constructed with DAS outcome data	0.09 (0.28)
Sick			0.10 (0.30)
Healthy			0.79 (0.41)
Inventory	Avg. inventory for the week of final outcome	Constructed with DAS outcome data and DAS daily report cards	425.93 (84.59)
Fostered	=1 if a dog was fostered	DAS outcome data	0.05 (0.22)
Foster returned	=1 if fostered dog was returned	DAS outcome data	0.05 (0.21)
Shelter time	# days in the shelter	DAS outcome data	7.47 (10.39)
Positive outcome	=1 if dog was adopted, transferred, fostered, or returned to the owner	DAS outcome data	0.75 (0.43)

Note: The number of intake records is n=259,228, the number of records with final outcomes recorded was 245,728. Remaining 13,500 records were missing final outcome.

3. The Analytical Framework

Our analytical framework integrates queuing theory and survival analysis to optimize grace period allocation across heterogeneous dog populations. Figure 1 provides an overview of the model components and their interactions. We discuss each component below.

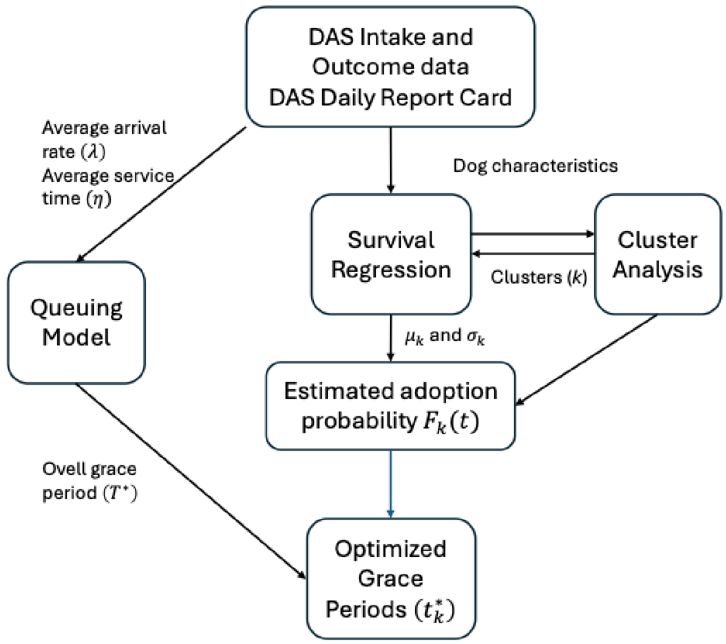

Figure 1. Overview of the analytical model.

3.1. Queuing Model

We model the shelter system as an $M/G/\infty$ queue with maximum time T, which captures municipal shelters' inability to reject arrivals. This structure assumes Poisson arrivals at rate λ (estimated from daily intake data), general service times with rate η (estimated via Kaplan-Meier methods to handle censoring, see on-line appendix B.1 for details), and unlimited servers with no blocking—reflecting that shelters must accept all animals.

Let S be an exponential random variable with rate η. We define a random variable $X_T = min\ \{S, T\}$. The variable X_T is the time in the system for a dog when the maximum length of stay is imposed as T.

A key advantage of the $M/G/\infty$ model is its mathematical tractability, even when the service time distribution for X_T is not purely exponential. This property ensures analytical simplicity and allows for closed-form calculations, as highlighted by Shortle et al. (2018, Section 6.2).

The expected number of dogs in the system under the grace period *T* is:

$$L(T) = \lambda \bar{x}_T \tag{3}$$

where $\bar{x}_T$ is the mean of variable X_T and we have $\bar{x}_T = (1 - e^{-\eta T})\frac{1}{\eta}$. Full derivations and mathematical assumptions for the queuing model are provided in on-line appendix B.2.

Following shelter industry standards and management input, we constrain the system such that shelter occupancy exceeds $(1 + m)$ times capacity at most τ percent of the time. For a shelter with *K* kennels, this requires finding the grace period T^* where:

$$P(N(T^*) \geq (1 + m)K) = \tau \tag{4}$$

where $N(T^*)$ follows a Poisson distribution with mean $L(T^*)$. Based on DAS practices and constraints, we operationalize m at 0.5, corresponding to 50% of kennels, and τ at 0.05, corresponding to 5% of the time hitting that capacity constraint.

Since the live release rate $LLR(T) = 1 - e^{-\eta T}$ is strictly increasing in *T*, maximizing LLR subject to the capacity constraint reduces to finding T^* such that a Poisson random variable with mean $L(T^*)$ exceeds $1.5K$ with probability 0.05.

3.2. Discriminant Analysis

The queuing model in Section 3.1 yielded a single aggregate service time, but adoption prospects vary widely across shelter dogs. To incorporate this heterogeneity while maintaining operational feasibility, we developed a three-step discriminant approach: regression, factor analysis, and clustering.

In Step 1, we used logistic regression to predict each dog's probability of euthanasia based on intake characteristics such as breed, age, health condition, and appearance (see Table 1). We excluded critically ill and fostered dogs from this analysis. The resulting predicted probability served as a summary measure of vulnerability. Including this probability in the clustering process ensures that resulting groups are meaningfully differentiated by survival outcomes.

Step 2 involved applying factor analysis to reduce the dimensionality of correlated intake variables. We retained four factors using standard criteria (eigenvalues > 1 and marginal improvements > 0.1), capturing key variance while eliminating redundancy. The factors were not rotated to preserve the raw statistical structure, and all factor scores were standardized to ensure comparability in the clustering stage.

In Step 3, we used k-means clustering on the standardized factor scores and predicted euthanasia probabilities to group dogs with similar characteristics and outcome trajectories. We tested multiple cluster solutions and selected five clusters based on the Calinski-Harabasz criterion, which balances separation between clusters with cohesion within clusters. We then reintroduced the excluded critically ill and foster groups as two additional categories.

3.3. Survival Analysis

To optimize grace periods across clusters, we estimate cluster-specific adoption probabilities $F_k(t_k)$ using a streamlined survival model for each year and cluster separately. We found that shelter performance varies greatly by year, partly due to managerial changes, partly due to COVID19, and likely also due to factors that we cannot observe. We are aware, for example, that in 2016 the city of Dallas commissioned a major study that stimulated improved performance. There may have been other similar events that affected shelter performance.

We employ two complementary approaches: (1) **Kaplan-Meier estimation** (nonparametric): Provides censoring-adjusted mean adoption times by year as inputs to the queuing model. (2) **Log-normal survival models** (parametric): Estimates smooth, differentiable adoption curves for each cluster and year. For each cluster k, we model:

$$F_k(t_k) = \Phi\left(\frac{ln(t_k) - \mu_k}{\sigma_k}\right) \tag{5}$$

where $\Phi(\cdot)$ is the standard normal CDF, and μ_k, σ_k are estimated via maximum likelihood using dog characteristics (size, age, health) and shelter conditions. These survival curves provide cluster-specific adoption trajectories essential for the marginal analysis in Section 3.4, where we allocate time based on $\frac{\partial F_k(t_k)}{\partial t_k}$.

3.4. Optimizing Grace Periods

Given the system-wide grace period T^* from the queuing model and cluster-specific adoption probabilities $F_k(t_k)$ from survival analysis, we now allocate time across clusters to maximize total positive outcomes.

We formulate the following nonlinear program:

$$\max \textstyle\sum_k w_k F_k(t_k) \tag{6}$$

Subject to

$$\sum_k w_k t_k \leq T^* - w_c t_c - w_F t_F \qquad \forall\, k$$

$$t_k \geq 0 \qquad \forall\, k$$

where w_k is the proportion of dogs in the cluster $k \in \{1 \dots 5\}$ w_F, and w_c are the proportions of dogs in the fostered and critical clusters, $F_k(t_k)$ is the cumulative adoption probability for the cluster k, evaluated at time t_k, t_F and t_c are the average times for critical and fostered dogs (fixed at empirical values).

We exclude critical dogs $F_c(t_c) = 0$ and fostered dogs $F_F(t_F) \approx 1$ from optimization as their outcomes are essentially predetermined. The constraint ensures the weighted average time across all clusters, including critical and fostered, equals T^*.

The optimization equalizes marginal adoption benefits across clusters at the optimum, implementing the Equi-marginal principle: $\frac{\partial F_k(t_k)}{\partial t_k}$ should be equal across all optimized clusters k at the solution. We solve using Sequential Least Squares Programming (SLSQP) given the smooth objective function and linear constraints.

Building on the system-wide grace period T^* identified in the queueing model (Section 3.1) and the adoption trajectories estimated through survival analysis (Section 3.3), we now refine the policy by allocating T^* across dog clusters. The goal is to maximize system-wide live release while respecting the capacity-implied limit on average time, by choosing a set of cluster-specific grace periods, $\{t_1, t_2, \dots, t_5\}$, each assigned to a distinct cluster based on its survival function $F_k(t_k)$.

This framework allows for differentiated treatment across dog groups. Instead of assigning all dogs the same maximum length of stay, we let adoption responsiveness—and not just population share—drive time allocations.

This optimization framework enables shelters to address adoption inequities through efficient resource allocation rather than simply extending stays. By equalizing marginal returns across dog types within capacity constraints, the approach recognizes that equity requires differentiated treatment based on need and potential benefit. The results demonstrate that effective shelter management balances compassion with operational reality: maximizing total adoptions often requires treating different dogs differently, making data-driven allocation both humane and sustainable.

3.5. Counterfactual Analysis

To evaluate our optimized grace period policy, we conduct counterfactual analysis from 2013-2025, comparing actual shelter performance against two policy alternatives using our cluster-based survival models. Under the Uniform Grace Period Policy, we calculate the live release rate achievable if all dog clusters received the same grace period T^*, set to the system-wide optimal value from our queuing model (Section 3.1). This represents performance under a single time limit applied uniformly across all dog types.

Under the Cluster-Optimized Policy, we apply our cluster-specific optimization framework (Section 3.4) to determine differentiated grace periods t^* for each dog type based on their adoption dynamics and marginal returns. The counterfactual analysis compares the actual LLR the shelter experienced to what could have been possible had our model been applied.

4. Results

4.1. Queuing Model Results

We use the queuing model to determine the system-wide optimal grace period T^* that maximizes live release while respecting capacity constraints.

Since the live release rate is defined as $LLR(T) = 1 - e^{-\eta T}$, and is strictly increasing in T, optimizing subject to a capacity constraint reduces to finding the grace period T^* such that the probability of exceeding critical capacity (1.5 times kennel space) is no more than 5%. Figure 2 illustrates this tradeoff for the year 2024, showing how LLR, average capacity, and the probability of exceeding 1.5 capacity evolve with the grace period. The x-axis plots the grace period T, and the y-axis ranges from 0 to 1.6, representing both kennel utilization and LLR (bounded between 0 and 1). The curve labeled P(Capacity > 1.5) tracks the probability that crowding exceeds the 1.5 capacity threshold; this threshold is operationally meaningful, as average occupancy approaches 1.4 when 561 dogs are housed in a 400-kennel facility. The dashed vertical line in the figure marks the optimal value $T^* = 23.596$ (also highlighted in Table 2), at which P(Capacity > 1.5) = 0.05. At this point, we also observe the maximum feasible live release rate LLR(T^*), representing the best outcome achievable under the binding capacity constraint, assuming a uniform grace period across all dogs.

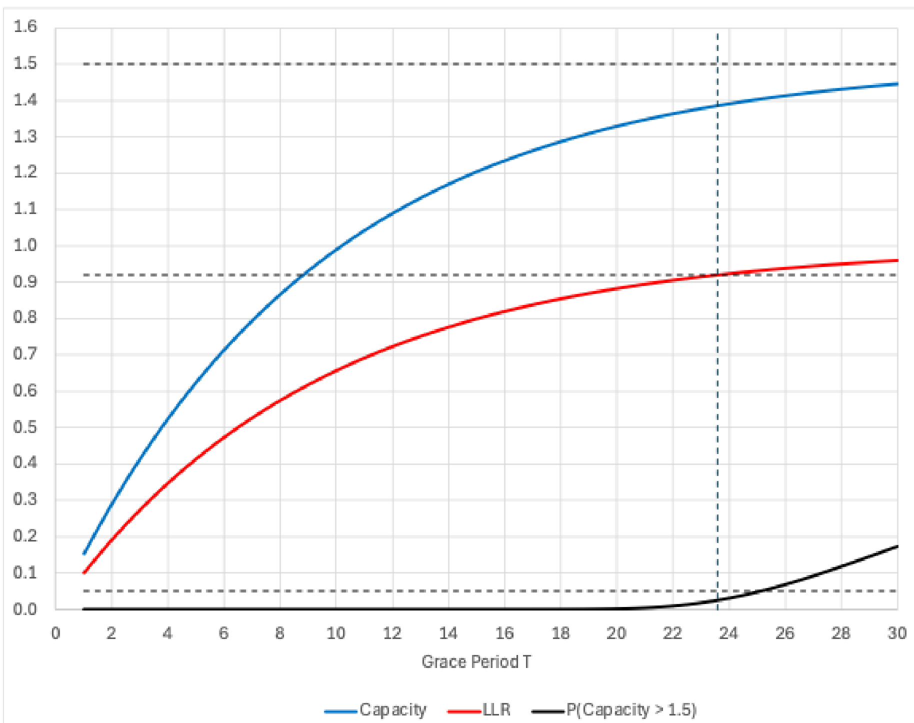

Figure 2. Determination of Optimal Grace Period T^* (Year 2024).

Because the shelter manages different number of animals in different years, we use the queuing model to optimize the average annual grace period for each year. Table 2 shows the inputs and outputs for each year. Year 2024 is shaded to show correspondence with Figure 2. Several patterns emerge. Years with low arrival rates relative to service time can sustain longer grace periods—for example, 2021 shows $T^* = 25.013$ days despite moderate service times because arrivals averaged only 36.87 dogs per day. Conversely, high-volume years such as 2013–2016 require shorter grace periods (14–18 days) to avoid overcrowding, even when service times are moderate.

The model's LLR predictions generally exceed observed values, suggesting that the model has the potential to reduce euthanasia. This gap is especially large in early years—e.g., 2013's actual LLR of 0.446 compared to the model's 0.579. The difference narrows in more recent years, with 2024 showing actual performance (0.867) slightly below the model (0.920), indicating possible adoption of policies influenced by our findings. We cap the model's LLR at 0.920 to reflect structural constraints, recognizing that roughly 8% dogs are too critically sick or injured to survive regardless of policy.

Table 2. Queuing model inputs and results by year compared to the observed LLR.

	Model Inputs			Queuing Results		Actual
Year	Mean Censored Days to Adoption ($1/\eta$)	Mean service rate (η)	Arrival Rate λ	T^*	Queuing LLR	LLR
2013	16.825	0.059	57.622	14.542	0.579	0.446
2014	14.814	0.068	57.748	15.798	0.656	0.483
2015	14.034	0.071	55.688	17.756	0.718	0.561
2016	13.044	0.077	57.639	17.884	0.746	0.644
2017	7.997	0.125	69.540	20.199	0.920	0.774
2018	8.430	0.119	76.132	17.469	0.874	0.813
2019	6.429	0.156	83.825	16.237	0.920	0.892
2020	5.417	0.185	40.208	13.682	0.920	0.926
2021	9.903	0.101	36.874	25.013	0.920	0.890
2022	13.333	0.075	43.816	33.676	0.920	0.791
2023	12.269	0.082	51.663	26.541	0.885	0.782
2024	9.342	0.107	64.464	23.596	0.920	0.867
2025	9.386	0.107	55.200	23.707	0.920	0.889

A sensitivity analysis is provided in on-line appendix C, showing that the optimal grace period T^* responds asymmetrically to changes in arrival rates and service times. When adoption is fast (low mean days to adoption), T^* remains relatively stable across a wide range of arrival rates. However, when adoption is slow (high mean days to adoption), even modest increases in arrival rates can sharply reduce T^*—for example, a 20% increase in intake may require cutting the grace period nearly in half. This sensitivity highlights that shelters facing slower adoption markets have less operational flexibility and must adjust grace periods more aggressively in response to intake surges.

4.2. Discriminant Analysis Results

We identified each cluster's primary characteristic, which turned out to be medical needs (non-critical), size, age under 6 months, and the number of previous shelter visits. Also, as mentioned earlier, we have the two manually constructed clusters: critical medical condition, and fostered.

Table 3. Summary of statistical clusters.

Dominant Characteristic	*Description*	*N*	*Avg. # days in Shelter*	*LLR*
Large and Medium	Large and medium dogs with moderate service times	88,307	9.39	0.71
Repeated Visitors	Dogs that visited the shelter at least two times, with more dogs returned to the owner	14,245	7.22	0.96
Small	Small healthy dogs with rapid positive outcomes	42,389	5.33	0.91
Baby	Puppies under 6 months old	37,717	6.66	0.87
Medical	Dogs requiring medical treatment	30,124	9.96	0.75
Fostered	Previously fostered dogs with near-certain positive outcomes	24,242	7.83	0.95
Critical	Critical cases with immediate outcomes	22,204	1.74	0.00

Note: Fostered dogs may have multiple records, but we consolidate them into a single observation when computing the actual LLR to align with industry practice on this table and throughout the paper.

Table 3 displays statistical clusters, their size, average time dogs in each cluster spend in the shelter, and the average *LLR* for each cluster. So, to provide context, the optimization model (section 3.4) reallocates average shelter time among the first five clusters in Table 3 while keeping the shelter time for the critical and the fostered clusters constant.

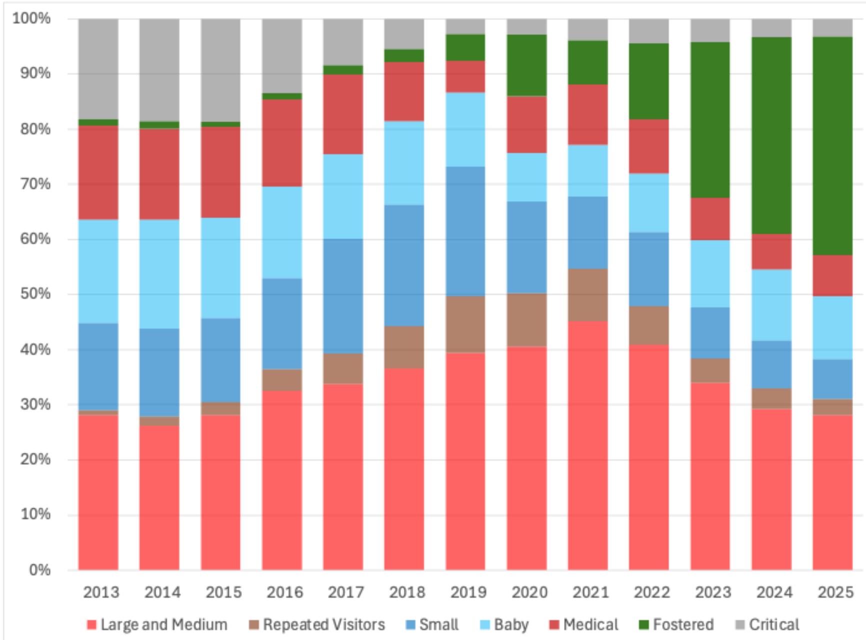

Figure 3. Clusters by year.

Figure 3 shows how the proportion of clusters changed over the years. The main pattern we would like to emphasize is that it appears that the fostering process was radically reformed around 2023.

The classification system we described above forms the foundation for the survival analysis and optimization that follow, where we will examine how each cluster's distinct adoption dynamics drive optimal time allocation.

4.3. Survival Analysis Results

Table 4 presents survival regression results for the complete dataset (years 2013-2025). Euthanized dogs are treated as right-censored observations. This analysis demonstrates substantial heterogeneity in adoption dynamics. The time ratios and their and 95% confidence intervals indicate how each characteristic affects expected service duration relative to baseline (medium-sized, male adult healthy dogs that are not of the breeds listed in Table 4, never fostered, and first time in the shelter).

Table 4. Survival Regression Results for years 2013-2025

Variables		Coefficient (SE)	Time Ratio	95% CI
Number of times a same dog visited the shelter	Number of visits	-0.242***(0.006)	0.785	[0.776, 0.794]
Size variables	Toy	-0.545***(0.019)	0.580	[0.558, 0.602]
	Small	-0.551***(0.010)	0.577	[0.566, 0.588]
	Large	-0.022***(0.008)	0.978	[0.963, 0.994]
	X-large	-0.229***(0.024)	0.795	[0.758, 0.834]
Gender	Female	0.015**(0.006)	1.015	[1.002, 1.028]
Primary breed includes	Bully	0.393***(0.009)	1.481	[1.455, 1.509]
	Husky	-0.375***(0.022)	0.687	[0.658, 0.718]
	Lab	0.186***(0.011)	1.204	[1.177, 1.232]
	Shepherd	0.019*(0.011)	1.019	[0.997, 1.042]
	Chihuahua	-0.097***(0.012)	0.907	[0.886, 0.929]
Appearance/Color	Black	-0.013*(0.007)	0.987	[0.973, 1.002]
	White	-0.103***(0.009)	0.902	[0.887, 0.918]
	Fluffy	-0.225***(0.016)	0.799	[0.774, 0.824]
Age of a dog at intake	Age	-0.027***(0.002)	0.973	[0.971, 0.976]
Younger than 6 months	Baby	-0.369***(0.009)	0.692	[0.680, 0.703]
Owner surrendered	Surrendered	0.199***(0.007)	1.220	[1.202, 1.238]
Health condition	Injured	0.715***(0.013)	2.045	[1.991, 2.099]
Health condition	Sick	0.590***(0.011)	1.805	[1.766, 1.845]
Shelter Load	Inventory	0.001***(0.00004)	1.001	[1.001, 1.001]
Ever fostered	Fostered	-1.500***(0.015)	0.223	[0.217, 0.230]
	Constant	1.642***(0.020)	5.165	[4.964, 5.374]
N	210601			
LL	-275398.40			
Chi-Squared (21)	39783.18 (21)			
μ	1.727	Median/Mean days	5.622/ 13.251	
σ	1.310			

Note: p-value < 0.01 ~ "***", < 0.05 ~ "**", < 0.1 ~ "*". We an individual dog as the unit of analysis to ensure that observations are independent, retaining only the most recent outcome for the dogs who were admitted multiple times.

Several patterns emerge that help validate our clustering approach. Medical conditions substantially extend service times, with sick (1.805) and injured (2.045) dogs requiring 81-105% more time. Small dogs achieve positive outcomes 42.3% faster than medium dogs (time ratio = 0.577), while large dogs show similar timing to baseline (0.978). Bully breeds require 48% longer service times (1.481), the largest effect among breed categories. Previous fostering accelerates positive outcomes by 78% (0.223), confirming the effectiveness of foster programs.

We selected the log-normal specification based on systematic comparison with alternative distributions (exponential, Weibull, loglogistic) using AIC and likelihood criteria that we present in the online appendix D.

Findings in Table 4 validate the key differentiators underlying our classification: medical status, size, age, and the number of visits, each significantly affect adoption timing in ways that justify separate policy treatment.

To demonstrate how adoption dynamics vary by cluster, we show in Figure 4 cumulative adoption probabilities aggregated over years 2013-2025, that we estimated separately for each cluster using lognormal specifications (critical cluster excluded). These curves reveal significant differences in how additional shelter time translates to positive outcomes. It can be seen from the figure that the marginal value of additional time—$\partial F_k(t_k)/\partial t_k$—varies dramatically across clusters.

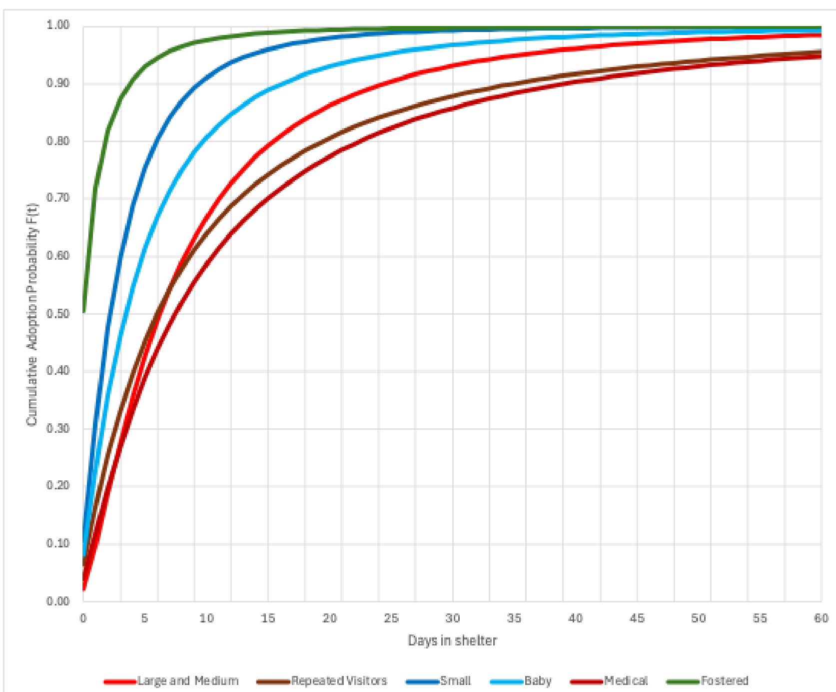

Note: Estimates are based on log-normal survival models.
Figure 4. Cumulative Adoption Probability $F_k(t_k)$ by Cluster (Years 2013-2025)

Fostered dogs show rapid success, reaching 90% by day 5. Small dogs achieve an 80% adoption probability by day 7 and plateau near 95% by day 14. Babies show similar adoption probability to Small dogs, reaching 80% by day 11 and plateauing near 95% by day 24. Large and Medium follow a middle trajectory, reaching 80% by day 17. Repeated visitors show slower but sustained gains, requiring 20 days to reach 80%. Medical dogs exhibit a gradual progression, reflecting extended treatment and recovery periods. This variation drives our optimization results.

4.4. Optimization Results

While cluster proportions shifted over time (particularly with foster program expansion), the fundamental adoption dynamics within each cluster remained relatively stable. We estimate parameters separately for each year/cluster combination and those estimates feed directly into the optimization model, where differences in adoption trajectories determine optimal grace period allocation. The persistent heterogeneity across clusters, combined with their varying marginal returns to time, creates opportunities for improving system-wide outcomes through strategic reallocation. We report parameter estimates for each cluster and year in the on-line appendix E.

Table 5 illustrates the outcome of our optimization approach for the year 2022. The other years are similar. Table 5 shows estimated log-normal parameters, cluster shares, and the resulting optimal grace periods t_k^* for each cluster. It also shows $LLR(T^*) = F_k(T^*) = \Phi\left(\frac{ln(T^*)-\mu_k}{\sigma_k}\right)$ which is what the LLR would have been for $t_k = T^* \; \forall \; k \in \{1,2,3,4,5\}$, and the cluster-specific $LLR^* = F(t^*) = \Phi\left(\frac{ln(t^*)-\mu_k}{\sigma_k}\right)$. The last column shows observed average service times calculated using the Kaplan-Meyer procedure to account for censoring.

Table 5. Optimal Allocation of Grace Periods by Cluster (Year 2022)

Cluster	μ	σ	Proportion w_k	Grace periods (t^*)	$LLR(T^*) = F(T^*)$	$LLR^* = F(t^*)$	Observed service time
Large and Medium	2.1861	1.0723	0.409	44.647	0.893	0.934	17.863
Repeated Visitors	2.3271	1.3472	0.070	52.617	0.811	0.888	10.238
Small	1.3016	0.9640	0.134	23.705	0.989	0.973	5.518
Baby	1.2329	1.0134	0.106	23.715	0.988	0.972	6.119
Medical	2.0807	1.1461	0.099	43.517	0.895	0.930	15.754
Fostered			0.137	12.170	0.928	0.928	12.170
Critical			0.045	1.418	0.000	0.000	1.418
Weighted Average				33.676	0.875	0.897	12.700

Note: Overall $T^* = 33.676$ based on the queuing model and overall observed LLR = 0.791 (Table 2).

For medium and large dogs, the actual LLR was 0.737 and the model would have increased it to 0.934, (marked with "+") while slightly decreasing LLR for repeated visitors (marked with "-"). This solution demonstrates the core logic of strategic time allocation: dogs with higher marginal adoption probabilities from additional shelter days receive longer grace periods. Despite Large and Medium dogs having lower baseline adoption rates (89.3% at uniform T^*), they receive the second to the longest grace periods (45 days) because their marginal returns remain high at longer durations. Conversely, Small dogs achieve 99% success at uniform allocation, leaving little room for improvement, so their time is reduced to 23.7 days. This reallocation improves system-wide LLR from 79.1% to 89.7%—seemingly modest but representing 1,559 additional dogs achieving positive outcomes in 2022 alone (see Table 6).

Figure 5 illustrates the mechanism driving optimal time allocation by plotting marginal adoption probabilities for Large and Medium dogs compared to the Small dogs for 2022. Small dogs show high marginal adoption probabilities in the first few days, but these gains decline rapidly as the most adoptable small dogs find homes quickly. Large and Medium dogs exhibit lower initial marginal probabilities but maintain more sustained gains over time, reflecting their need for extended exposure to find suitable adopters. This divergence explains why reallocating time from Small to Large and Medium dogs increases total adoptions: the marginal benefit of additional days is higher for Large and Medium dogs beyond the initial period, even though Small dogs have higher baseline adoption rates. Tradeoff between these two clusters in other years is similar.

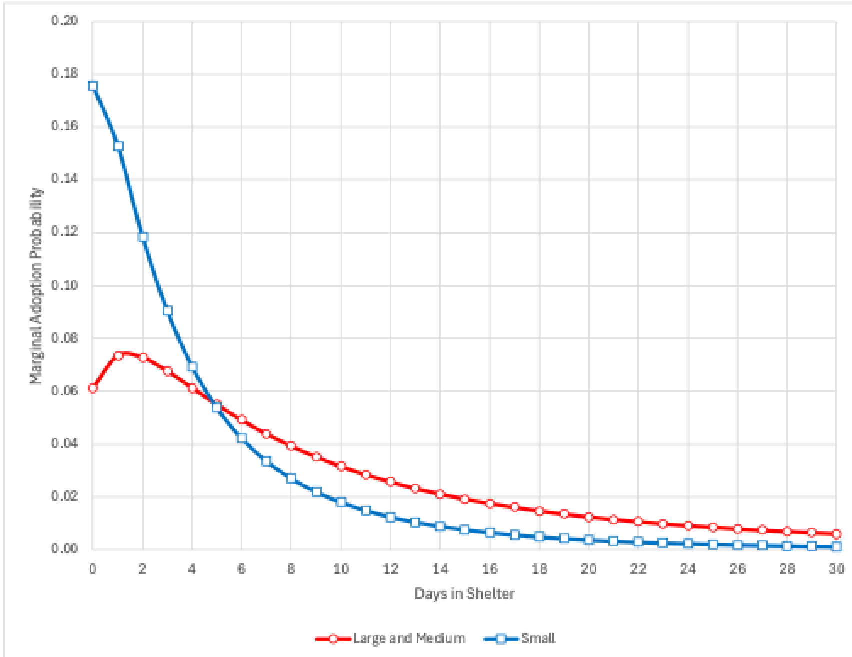

Figure 5. Tradeoff between Large and Medium and Small Allocations (Year 2022)

This analysis reveals a fundamental insight: uniform grace periods systematically misallocate scarce shelter resources. Strategic reallocation based on adoption dynamics generates substantial improvements by directing time toward dogs with the highest marginal returns. Large and Medium dogs exemplify this principle—their sustained adoption gains over time create high value when allocated additional days, while reallocating time away from dogs with diminishing returns (Small dogs after the initial period) improves system-wide outcomes.

5. Policy Evaluation and Counterfactual Analysis

Table 6 compares historical performance with both counterfactual policies across the study period. Results vary substantially by year, reflecting how operational conditions affect optimization potential. In poorly performing years (2013-2016), the model suggests saving 1,000-2,000 additional dogs annually. During the COVID years (2020–2021), when historical performance was already exceptionally high, our model still identifies scope for further LLR improvement, corresponding to an additional 100–400 dogs saved. In more recent years (2022–2024), as shelter management practices improved, the model continues to reveal opportunities to save 1,000–1,700 additional dogs annually, underscoring its value even under favorable operational conditions. Overall, our model has the potential to save an average of 1,434 dogs per year over the 2013–2025 period.

Table 6. Historical and Counterfactual Live Release Rates (2013–2025)

Year	N	Positive	Actual LLR	Homogeneous T^*	$LLR(T^*) = F(T^*)$	$LLR^* = F(t^*)$	Dogs Saved
2013	20,841	9,288	0.446	14.542	0.466	0.524	1,643
2014	20,965	10,132	0.483	15.798	0.514	0.577	1,965
2015	20,261	11,373	0.561	17.756	0.599	0.654	1,872
2016	20,978	13,513	0.644	17.884	0.652	0.694	1,044
2017	25,058	19,397	0.774	20.199	0.844	0.864	2,253
2018	27,463	22,329	0.813	17.469	0.871	0.892	2,159
2019	29,925	26,708	0.892	16.237	0.933	0.948	1,663
2020	14,048	13,004	0.926	13.682	0.923	0.936	143
2021	12,620	11,237	0.890	25.013	0.906	0.919	361
2022	14,661	11,590	0.791	33.676	0.875	0.897	1,559
2023	15,800	12,362	0.782	26.541	0.848	0.887	1,653
2024	19,014	16,481	0.867	23.596	0.890	0.923	1,060
2025	4,094	3,640	0.889	23.707	0.906	0.936	192

Note: Annual average is 1,434 dogs saved per year. In 2025, we have 0.25 year data.

We plot these results graphically in Figure 6, which compares the actual historical LLR, the counterfactual based on a uniform grace period, and the cluster-specific grace period. Both counterfactual policies usually outperform actual outcomes, particularly in earlier years when shelter decisions deviated most from model recommendations. The performance gap narrows in recent years (2023-2025) as actual practices improved, potentially reflecting increased awareness of the benefits of differentiated grace periods based on numerous discussions we had with the shelter management, even if implementation methods differed from our specific framework. Even under good shelter operations, our model projects annual savings of over 1,000 dogs, while poorly performing periods show potential gains exceeding 2,200 dogs annually.

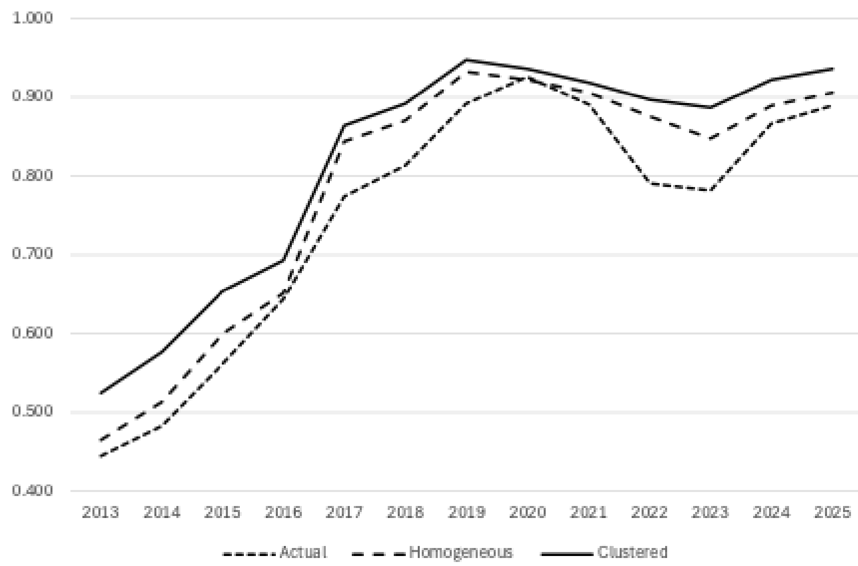

Figure 6. Comparison of Historical and Counterfactual LLRs (2013–2025)

This modeling framework extends beyond Dallas Animal Services to any municipal shelter with open-admission policies and basic data infrastructure. Implementation requires only daily arrival counts, censored length-of-stay records (historical), and dog-level attributes for clustering and survival estimation—data routinely collected by most shelters. This approach therefore provides a scalable, data-driven framework for humane decision-making across diverse shelter environments.

6. Discussion and Conclusion

This research demonstrates that strategic time allocation can substantially improve animal shelter outcomes without additional resources. By treating time as an optimization variable rather than a constraint, shelters can achieve more humane and efficient operations even under scarcity. Our counterfactual analysis suggests potential annual savings of more than 1,400 dogs, even when a shelter is well-managed, with gains exceeding 2,200 dogs during periods of poor performance.

Our analysis reveals that even well-established operational practices can benefit from systematic optimization. The finding that slower-adopting dogs should receive longer grace periods challenges conventional wisdom that might prioritize easily-adoptable dogs. This insight emerges from recognizing diminishing marginal returns—while faster-adopting dogs succeed quickly, extended time for challenging cases can yield substantial aggregate improvements. While emotionally we might prefer saving dogs that need minimal help, the mathematics favor allocating scarce time where it makes the most difference. A small dog will likely find a home regardless; a large or medium dog's fate may hinge on those extra days of exposure.

This framework’s contribution to operations research methodology is in demonstrating how survival analysis can be integrated with queuing models to capture heterogeneity in service systems, thereby offering a generalizable template for resource allocation under scarcity. The approach of applying queuing models with heterogeneity adjustments has wide applicability wherever limited resources must be allocated across clients with diverse needs. Potential applications extend well beyond animal shelters. In homeless shelters and transitional housing, exit timing depends on both individual circumstances and the availability of permanent housing (Fisher et al., 2014). In workforce training programs, participants require different levels of support before becoming job-ready (Rosenheck and Mares, 2007). In healthcare settings, the framework is particularly promising: emergency department patient flow varies by treatment complexity and discharge timing; outpatient discharge planning must adapt to individual recovery trajectories; and substance abuse treatment facilities face highly variable patient needs (Green et al., 2006; Duguay and Chetouane, 2007).

The framework’s contribution to animal welfare practice lies in providing an implementable classification system that translates statistical clusters into clear operational decision rules. This enables shelters to assign grace periods in real time based on observable dog characteristics. Implementation requires only minimal infrastructure—daily intake counts, length-of-stay records, and basic descriptors already collected by most shelters. As a practical entry point, shelters can begin by tracking service times by dog type within our operational categories, calculating the current average grace period for each cluster, and gradually adjusting these values toward model recommendations while monitoring capacity. To support daily decisions, a simple spreadsheet-based tool can flag dogs approaching their cluster-specific grace period, ensuring euthanasia choices are made in a transparent, systematic, and data-driven manner.

This observational study faces several limitations. First, we cannot establish causal relationships between time allocation and outcomes—shelters may already give longer stays to dogs they perceive as more adoptable, biasing our survival estimates upward. Second, unobserved factors may confound results: staff advocacy for specific dogs, seasonal adoption patterns, or marketing efforts could influence outcomes independent of time allocation. Third, our treatment of fostered dogs as exogenous oversimplifies reality—foster placements likely select for dogs with better prospects, and the 95% success rate may not generalize if foster programs expand dramatically. Finally, our static optimization assumes stable conditions, while real shelters face dynamic shocks (disease outbreaks, hoarding cases, seasonal surges) that require adaptive responses.

Several extensions could enhance this framework. First, dynamic optimization could adjust grace periods in real-time as conditions change—intake surges, staffing fluctuations, or capacity shifts. Second, behavioral interventions deserve investigation: how do adopter outreach, volunteer programs, or facility improvements interact with optimized timing? Machine learning approaches might predict at-risk dogs at intake, enabling proactive grace period assignment and networking strategies. Understanding foster care dynamics represents another promising direction, particularly as foster-to-adopt programs expand and integration with shelter management systems improves. Finally, validation across diverse shelter contexts—smaller facilities, different intake policies, varying community demographics—would establish broader applicability.

The fundamental insight transcends animal welfare: when resources are constrained and clients have heterogeneous needs, differentiated treatment based on marginal returns can improve both efficiency and equity. This principle offers a rigorous foundation for data-driven management across service systems in which scarcity demands difficult allocation decisions.

7. References

Benade G, Alptekinoglu A (2024) Achieving Rawlsian justice in food rescue. SSRN Working Paper. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4992842

Best Friends Animal Society (2024) 2024 National data report. Retrieved on 06/10/2025 from https://bestfriends.org/network/studies-publications/shelter-pet-lifesaving-data-2024-report

Castle J (2024) 3.4% fewer animals were killed in shelters so far in 2024. *Best Friends Animal Society.* https://bestfriends.org/stories/julie-castle-blog/34-fewer-animals-were-killed-shelters-so-far-2024

Chehrazi N, Sanders RE, Stamatopoulos I (2025) Inventory information frictions explain price rigidity in perishable groceries. *Marketing Science. 44*(2), 411–436.

Duguay C, Chetouane F (2007) Modeling and improving emergency department systems using discrete event simulation. *Simulation*. 83(4), 311–320.

Fisher BW, Mayberry LS, Shinn M, Khadduri J (2014) Leaving homelessness behind: Housing decisions among families exiting shelter. *Housing Policy Debate*. 24(2), 364-386.

Gaur V, Honhon D (2006) Assortment planning and inventory decisions under a locational choice model. *Management Science*. 52(10), 1528–1543.

Green LV, Savin S, Wang B (2006) Managing patient service in a diagnostic medical facility. *Operations Research*. 54(1), 11–25.

Jalayer A, Jalayer M, Khakzand M, Faizi M (2024) Automated optimal layout generator for animal shelters: A framework based on genetic algorithm, TOPSIS and graph theory. *arXiv preprint.* arXiv:2405.14172.

Kang JH, Han J (2019) Optimizing the operation of animal shelters to minimize unnecessary euthanasia: A case study in the Seoul Capital Area. *Sustainability.* 11(23):6702. https://doi.org/10.3390/su11236702.

Li H, Shen Q, Bart Y (2021) Dynamic resource allocation on multi-category two-sided platforms. *Management Science*. 67(2), 984-1003.

Pinçe Ç, Gürler Ü, Berk E (2008) A continuous review replenishment–disposal policy for an inventory system with autonomous supply and fixed disposal costs. *European Journal of Operational Research. 190*(2), 421–442.

Rosenheck RA, Mares AS (2007) Implementation of supported employment for homeless veterans with psychiatric or addiction disorders: Two-year outcomes. *Psychiatric Services*. 58(3), 325-333.

Shelter Animals Count (2024) National database statistics. Accessed June 13, 2025 https://www.shelteranimalscount.org/explore-the-data/statistics-2024

Shortle JF, Thompson JM, Gross D, Harris CM (2018) *Fundamentals of queueing theory. (*John Wiley & Sons.)

The Association of Shelter Veterinarians (2022) The Guidelines for Standards of Care in Animal Shelters: 2nd edition. *J. Shelter Medicine and Community Animal Health.* 1(S1):1–76. https://doi.org/10.56771/ASVguidelines.2022.

Turken N, Carrillo JE, Paul A (2021) The impacts of mergers, capacity expansion and adoptions on animal shelter performance: A queuing analysis. *European Journal of Operational Research*. 292(1), 299-310.

University of Wisconsin–Madison Shelter Medicine Program (2023) Calculating Shelter Capacity. Published August 9, 2023. Accessed May 31, 2025. https://sheltermedicine.wisc.edu/library/resources/calculating-shelter-capacity

Weiss E, Patronek G, Slater M, Garrison L, Medicus K (2013) Community partnering as a tool for improving live release rate in animal shelters in the United States. *Journal of Applied Animal Welfare Science*. 16(3), 221-238.