Regress, Reverse, Recycle: Contextual Stochastic Optimization in Waste Policy and Logistics Network Design

Austin Saragih
Center for Transportation and Logistics, Massachusetts Institute of Technology, Cambridge, MA, USA, saragih@mit.edu

Milena Janjevic
Center for Transportation and Logistics, Massachusetts Institute of Technology, Cambridge, MA, USA, mjanjevi@mit.edu,

Yossi Sheffi
Center for Transportation and Logistics, Massachusetts Institute of Technology, Cambridge, MA, USA, sheffi@mit.edu,

Jan C. Fransoo
School of Economics and Management, Tilburg University, Netherlands jan.fransoo@tilburguniversity.edu

Effective policies and reverse logistics networks for Municipal Solid Waste (MSW) recycling are crucial for advancing the circular economy. Current approaches to MSW recycling often decouple reverse logistics from endogenous recycling policies and separate waste collection routing from network design, failing to capture critical interdependencies. We address these limitations by incorporating endogenous recycling policy estimation and collection routing into Reverse Logistics Network Design (RLND). Our methodology uses Post Double Selection with Rigorous Lasso (PDS RLasso) to regress recycling rates against municipal policies and characteristics, then optimizes the reverse logistics network using Empirical Residuals-based Sample Average Approximation (ER-SAA). This approach enables the transformation of endogenous policies into exogenous ones for optimization. Through a comprehensive case study of Massachusetts' recycling network, we provide data-driven, actionable policy options to improve recycling rates and profitability. This methodology extends beyond MSW applications to broader supply chain network design and stochastic programming problems.

Key words: waste policy; recycling; contextual stochastic optimization; reverse logistics network design

1. Introduction

Recycling is crucial to global environmental sustainability and resource conservation (Anshassi and Townsend 2023). The United States (US), a major contributor to global waste generation, faces significant recycling challenges, as less than a third of the 292.4M tons of Municipal Solid Waste (MSW) were recycled in 2018 (EPA 2020). Recognizing this issue, the US government has a 50% recycling rate goal by 2030 (EPA 2021), Massachussetts has 30% (MassDEP 2021).

Achieving higher recycling rates relies on several factors, such as improving waste separation, expanding the recycling infrastructure, promoting public participation, and enhancing policies that support the circular economy (Basuhi et al. 2021, 2024). These factors are closely tied to the effectiveness of the MSW recycling system, which depends on three interconnected levers: (1) policy measures, (2) strategic design of the reverse logistics network, and (3) tactical operation planning of recycling operations (Saragih et al. 2024a). Policy measures include regulations (e.g., mandatory recycling programs), economic levers (e.g., taxes and incentives), and public awareness campaigns, which are done by the authorities of the state and its municipalities. A key policy decision is the *choice of municipal waste collection strategy.* Specifically, municipalities must decide whether to implement curbside collection (i.e., pickup of recyclables directly from households), drop-off systems (i.e., systems where residents transport recyclables to a designated point), or no recycling policy. Each collection strategy has distinct cost and accessibility implications. While municipalities focus on policy decisions, Materials Recycling Facility (MRF) operators manage the *strategic design of the reverse logistics network* (i.e., choices around the placement of logistics facilities), as well as the *tactical operations planning* (i.e., optimizing collection routes from the households or drop-off points to the MRF facilities). Decisions in these three areas are highly interdependent. For example, municipal policies directly influence recycling volumes, which subsequently drive the reverse logistics network's capacity requirements and the routing operations' efficiency. Similarly, the reverse logistics network design affects the routing operations' efficiency. An integrated optimization of these levers is crucial for developing an efficient recycling system.

Current approaches in the literature typically consider these levers independently, failing to capture these interdependencies. Studies on recycling policies (e.g., Kinnaman (2006)) are mainly empirical and focus on the impact of collection strategies, landfill taxes, or information provision on recycling rates. They do not examine the relationships between policies and decisions relevant to strategic network design or operational planning. The strategic design of the reverse logistics network is typically addressed using deterministic optimization models for Reverse Logistics Network Design (RLND) (e.g., Fleischmann et al. (2001)). Although examples of stochastic models exist (e.g., Bing et al. (2014a)), they assume that recycling policies are predetermined and do not influence the uncertain scenarios of recyclable supply. This approach overlooks evidence in the literature indicating that changes to these policies significantly impact recycling rates and supply. Consequently, we need a more advanced modeling approach that captures how policy decisions influence uncertain recycling rates, as in endogenous Contextual Stochastic Optimization (CSO) models (Zhu et al. 2024). In addition, tactical and operational planning is explored mainly through research on collection routing (e.g., Bing et al. (2014b)), which is often disconnected from network design, as highlighted in Bing et al. (2016). These fragmented approaches limit our understanding of the interaction between different levers and ultimately lead to less effective recycling systems.

This paper resolves this problem by proposing an integrated approach combining policy, network design, and tactical and operational planning described in this study Section 2. Building on the literature review in Section 3, this paper advances three key contributions that extend beyond existing approaches:

1. Drawing on data and a theoretical framework for Massachusetts (MA) MSW system (Section 4), we model an *integrated optimization framework* that simultaneously determines municipal recycling policies and reverse logistics network design (Section 5). First, we model recycling rates as a stochastic outcome impacted by both policy decisions and municipal characteristics or covariates, resulting in an *endogenous CSO model*. Second, we integrate strategic network design decisions with tactical and operational planning by *extending the Continuum Approximation (CA) approach* to include routing cost approximations that reflect the endogenous recyclable supply.
2. Given that policy endogeneity renders the problem intractable (Zhu et al. 2024), we develop a *novel solution approach* to enable policy optimization under uncertainty (Section 6). This approach extends beyond existing studies that focus solely on exogenous uncertainties. We draw from Kannan et al. (2025), Belloni et al. (2014), and Belloni et al. (2012), to propose a method integrating Empirical Residuals-based Sample Average Approximation (ER-SAA) and Post Double Selection with Rigorous Lasso (PDS RLasso) that converts endogenous policies and uncertainties into exogenous ones (Proposition 1). Our approach also leverages contextual heterogeneity through municipal characteristics, enabling tailored policy decisions that reflect varying recycling rates across municipalities.
3. We apply our method to the Massachusetts MSW recycling network case study and data (Section 7), offering *data-driven analysis and policy options* that the state can implement to improve its recycling rate and be closer to its 2030 goal of 30% (MassDEP 2021). Specifically, we demonstrate that our approach can lead to an increase in the recycling rate by 7.5%, reduce the cost of public authorities by $20.1M (22%), and increase profitability to MRF operators from a loss of $5.2M (-4% of cost) to a profit of $7.5M (7% of cost) (Section 7).

Although we focus on recycling, our contributions have broader implications. Many supply chains involve decisions regarding the design of policies or services, network configuration, and operational planning under uncertainty. In e-commerce supply chains, demand is endogenous and depends on service offerings and customer preferences (Janjevic et al. 2021), similar to how recycling policies influence recycling volumes. There is also a strong interdependence between strategic network design and operational planning (Janjevic et al. 2021). With this in mind, our study's general methodology and solution approach can be applied to a broader set of network design and stochastic programming problems, leveraging tools from causal inference under conditional independence assumption.

2. Problem Setting

This section describes the general problem setting in two steps. Section 2.1 details the MSW recycling network, policies, and actors. Section 2.2 formally defines and formulates the strategic RLND problem under consideration.

2.1. MSW Recycling Network, Policies and Actors

This study focuses on residential MSW recycling network and policies. Here, we also focus exclusively on recyclable materials since non-recyclable waste is sent directly to landfills. Therefore, commercial and industrial waste streams are outside the scope of this study.

Figure 1 presents a high-level overview of the reverse logistics network. Residents in municipalities generate MSW recyclables, which are collected through two mechanisms. In *curbside collection*, residents place recyclables in designated bins for scheduled pickup. In *drop-off collection*, residents take recyclables to designated drop-off points, where MRF operators coordinate regular line-haul pickups. Recyclable materials are sorted and sold at the MRFs, where residuals may be landfilled or incinerated. Though relevant economically and environmentally, the processes beyond MRFs lie outside our primary scope.

Figure 1 **Illustration of residential MSW recycling and collection methods: curbside and drop-off**

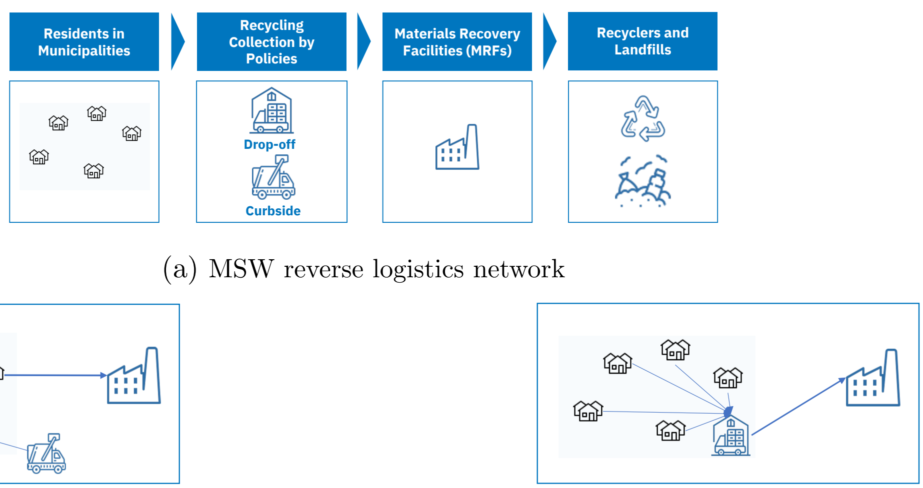

(a) MSW reverse logistics network

(b) Curbside collection

(c) Drop-off collection

This system represents a complex interaction between several stakeholder groups.

Residents are the generators of recyclable materials. Their recycling habits directly impact recycling rates and volumes. Their behavior is shaped by convenience—curbside collection being the more convenient option—and other factors, such as their environmental concern.

Public authorities are responsible for defining recycling policies, i.e., choosing between the curbside pickup and drop-off systems and working with MRF operators for their implementation through public-private partnerships. Their objective is to minimize recycling operations' total direct and indirect costs. Direct costs include infrastructure investments (i.e., bins for curbside programs and maintenance of drop-off facilities) and tipping fees paid by public authorities to MRF operators for collecting and processing recyclables. Indirect costs include environmental and social costs associated with disposal, incineration, and indirect loss of economic value.

MRF operators manage the design of the reverse logistics network and the daily operations. Their primary objective is to maintain profitability. They optimize facility locations and material flows from the collection of recyclables, which are then processed at their facilities. Typically, curbside collection is more costly than drop-off collection.

Decisions among stakeholders, public authorities, residents, and MRF operators are deeply interconnected. The choice of curbside and drop-off collection by authorities shapes the recycling behavior of residents, affecting the volume of recyclables and the efficiency of collection, which in turn impacts the profitability of MRF operators and the costs of the network. However, recycling outcomes are not solely determined by these interactions; external factors that explain heterogeneity, such as municipal demographics, socioeconomic characteristics, and environmental attitudes, also play a significant role. Thus, a comprehensive framework is required to account for both *endogenous factors* (stakeholder interactions) and *exogenous factors* (external characteristics), ensuring effective and tailored recycling policies and system designs.

2.2. Contextual RLND with Endogenous Policies and Uncertainties

In this section, we formalize the Contextual Reverse Logistics Network Design (RLND) problem. The problem is characterized by two key elements, which are presented below.

Integrated policy, network design, and collection routing decisions. This first element is the integration of three aspects. The first aspect is *policy decisions*. Here, the objective is to decide where to implement recycling policies and to choose between curbside or drop-off collection. The second aspect is *network design decisions*. The objective is to select which MRFs to open and decide on the allocation of municipalities to MRFs while respecting the profitability and capacity constraints of the MRFs. The third aspect is *collection routing decisions*. Once recyclables are collected, they must be routed to the MRFs efficiently, requiring the inclusion of routing decisions in our model. Note that we consider a large-scale problem setting with millions of households served.

Endogenous Policies and Uncertainties. Our second element includes two key interdependencies: (1) policy decisions are influenced by local municipal characteristics such as population density, income, and education levels, and (2) recycling rate distributions depend on both policy decisions and these municipal characteristics. This interplay introduces endogeneity into the system, where policies and outcomes are mutually dependent. These interdependencies must be captured for a realistic and actionable framework to design efficient and sustainable municipal recycling systems.

In addition to this endogeneity, estimating recycling rates under different policies introduces uncertainty. Recycling rates cannot be predicted with complete accuracy when policies are changed. These uncertainties arise from two sources: reducible uncertainties that can be explained by municipal characteristics and irreducible uncertainty that remains even after accounting for these factors. Overall, these uncertainties and policies are endogenous.

3. Related Literature

We review five research areas to address the problem setting described in Section 2.2. First, our problem requires estimating the effect of recycling policies on recycling rates. To this end, we review empirical studies on MSW recycling policies (Section 3.1). Second, we consider a network design problem with an uncertain recycling supply. This problem aligns with the literature on reverse logistics network design (RLND) under uncertainty (Section 3.2). Third, operational and tactical routing costs are critical in system efficiency. Here, we first review Waste Collection Routing (WCR) models (Section 3.3), which focus on vehicle routing for recyclable collection. Given that exact vehicle-routing models become computationally intractable for large-scale systems, we review the continuum approximation (CA) methods (Section 3.3), which offer scalable estimates of routing costs. Finally, our problem requires combining predictions of recycling rates with network optimization. To this end, we review contextual stochastic optimization (CSO) methods (Section 3.4). We conclude by identifying gaps in current approaches in Section 3.5.

3.1. Empirical Research in MSW Recycling Policy

Studies in this area consistently demonstrate that policies and socioeconomic factors influence recycling rates. Specifically, the effectiveness of recycling policies is interrelated with recycling volumes, pricing structures, and community characteristics (Kinnaman and Fullerton 2000, Sidique et al. 2010). Curbside recycling programs are the primary driver of participation (Jenkins et al. 2003, Gellynck et al. 2011, Mueller 2013), offering convenience that encourages recycling behavior. Recent research highlights how information provision and penalties can also improve recycling performance (McKie et al. 2024), while drop-off centers remain vital alternatives where curbside collection is impractical (Folz 1995). Finally, although income levels correlate significantly with recycling rates (Gellynck et al. 2011, Park 2018), social and cultural factors also play a crucial role (Laidley 2013). A key insight from this literature is the complex interplay between endogenous and exogenous factors:

- *Endogenous policies:* Municipal recycling policies are endogenous, as they are influenced by municipal characteristics (Kinnaman and Fullerton 2000, Sidique et al. 2010).
- *Exogenous factors:* Community characteristics such as demographics, education levels, and cultural norms are largely exogenous, as they are determined outside the recycling system (Gellynck et al. 2011, Laidley 2013, Park 2018).

- *Endogenous recycling rate:* These are influenced by both endogenous policies and exogenous factors (Kinnaman and Fullerton 2000, Sidique et al. 2010), making them a complex outcome to predict and optimize.

Despite the wealth of descriptive analyses in these studies, there remains a significant gap in translating these insights into prescriptive solutions. Specifically, the literature lacks integrated approaches that combine empirical findings with optimization techniques to improve recycling network design and performance simultaneously. This gap presents an opportunity for research that bridges descriptive analyses with actionable and optimized policy recommendations.

3.2. Municipal Solid Waste Reverse Logistics Network Design Under Uncertainty

Fleischmann et al. (2001) first popularize the idea of designing a network of facilities for reverse logistics by developing generic RLND to enable a closed-loop supply chain. In the context of recycling, Demeester et al. (2013) develop a simple RLND model for recycling facilities and analyze the interaction of different scales in the formation of efficient recycling facility networks. These works focus on general reverse flows of goods, whereas our context is the reverse flows of MSW. To that end, we review the relevant literature for our MSW problem.

Bing et al. (2014a) formulate a stochastic MSW RLND model considering separation methods and the interests of stakeholders, leading to cost savings and a significant reduction in carbon emissions. Gambella et al. (2019) formulate a stochastic MSW RLND model and show the value of the stochastic solution compared to the average-case solution. More recently, Reddy et al. (2022) incorporate carbon emissions and vehicle selection into the RLND model and solved it with a revised form of the Benders decomposition algorithm. Saragih et al. (2024b) integrate RLND with policy for deposit refund systems of MSW. We refer the reader to reviews by Govindan et al. (2015) and Govindan et al. (2017) for general RLND and RLND under uncertainty, respectively. We also refer to Van Engeland et al. (2020) and Bing et al. (2016) for MSW RLND models. From this review, we have two observations. First, current studies focus primarily on facility location and flow decisions, treating recycling policies as predetermined rather than decision variables. Second, all models consider flows from consumers to facilities as simple network flows. In reality, these flows come from WCR operations, which collect recyclables from the curb.

3.3. Waste Collection Routing and Routing Continuum Approximation (CA) in Network Design

WCR research builds on foundational vehicle routing work of Beltrami and Bodin (1974) and Angelelli and Speranza (2002). Notable advances include Hemmelmayr et al. (2014)'s integration of routing with bin allocation and Bing et al. (2014b)'s application to MSW collection under uncertainty. Recent work by Spinelli et al. (2024) extends these models to multi-period stochastic scenarios. However, these routing models typically focus on operational and tactical planning and are disconnected from strategic network design decisions. Integrating such models with strategic network design decisions results in tractability challenges, as noted by Winkenbach et al. (2016).

To address this issue, a common approach in the literature is to approximately model the tactical and operational decisions through CA methods. CA methods allow bridging routing operations and network design by tractably approximating routing decisions. While no specific application of CA exists in MSW, we can note several works that apply this to a more general last-mile distribution setting. For example, Winkenbach et al. (2016) develops optimal routing cost estimation formulas for large-scale forward logistics network design problems (also known as supply chain network design) based on Daganzo (1984). Janjevic et al. (2019) extends this approach with extensive numerical validation, while Janjevic et al. (2021) demonstrates its effectiveness in complex, multi-echelon e-commerce distribution networks. We note that these works typically consider fixed input demand (supply in our case). A contextual form of routing CA is needed.

3.4. Contextual Stochastic Optimization

The integration of machine learning with optimization enables decision makers to take advantage of contextual or covariate data to improve decision making. This area of research has produced three main CSO frameworks, as reviewed by Sadana et al. (2024): sequential learning and optimization, integrated learning and optimization, and decision rules. Sequential approaches, such as those proposed by Bertsimas and Kallus (2020) and Kannan et al. (2025), first use machine learning to predict uncertain parameters, then optimize based on these predictions. Integrated approaches, such as by Elmachtoub and Grigas (2022), combine learning and optimization to minimize decision errors directly. Decision rule approaches focus on learning optimal decision policies directly (Ban and Rudin 2019). In this paper, we focus on sequential approaches since Elmachtoub et al. (2023) demonstrate that the sequential approach has superior performance compared to integrated learning and optimization and regular Sample Average Approximation (SAA) when the model is specified correctly.

ER-SAA, introduced by Kannan et al. (2025), has established itself as a key method within the sequential approach. It extends SAA (Kleywegt et al. 2002) by leveraging covariates to derive both the predictions and their corresponding residuals. These outputs, which come from a trained machine learning model, are then used to generate scenarios for stochastic programming. This method can be applied to stochastic mixed-integer linear programs (Kannan et al. 2025), unlike most of the current applications of CSO that have focused on linear problems, such as optimizing production and shipment (Bertsimas and Kallus 2020).

We want to point out that these studies have only considered exogenous settings, where the predicted values only depend on exogenous covariates. For example, in Kannan et al. (2025), the covariates of the customers can predict the uncertain demand of the customers. That is, the dependent variable only depends on independent covariates or contextual data. These could be features of the product/service or about the customers. However, in broader settings, including ours, the uncertain variables depend on covariates and the decision variables in the optimization problem. These problems are known as ER-SAA with endogenous uncertainty (Zhu et al. 2024). To handle endogenous uncertainty, Zhu et al. (2024) use a distributionally robust approach with ambiguity sets that can handle endogeneity instead of ER-SAA. For our study, instead, we choose to focus on extending ER-SAA to settings that can handle endogeneity.

3.5. Summary and Literature Gap

The literature reveals several disconnects between empirical policy studies, stochastic reverse logistics network design (RLND), and operational waste collection routing (WCR) models. First, while empirical research establishes the importance of recycling policies (particularly curbside versus drop-off) and demographic factors, most of these studies only *describe* policy impacts rather than *prescribing* optimal policies and recycling networks. Our study guides the transformation of descriptive empirical studies into data-driven prescriptive analytics. Second, the majority of MSW RLND work treats policy decisions as exogenous and focuses on facility location and flow optimization, overlooking how municipal policies dynamically influence supply (and thus facility utilization). Third, vehicle routing research in WCR focuses on operational models, typically decoupling the routing lever from strategic network design and policy factors. In addition, in large-scale networks with many municipalities and households, exact routing models are computationally intractable. Continuum approximation (CA) offers a tractable way to model routing costs at scale, but prior CA applications do not address contextual recycling volumes.

Moreover, although recent CSO methods can integrate machine learning predictions with optimization, they are rarely applied in an *endogenous* setting, where the decision itself (policy) influences the uncertain parameter (recycling rate). This creates additional complexity not addressed by existing ER-SAA-based approaches.

4. Recycling Rate Analysis - Data and Theoretical Framework

This section presents the data and theoretical framework (i.e., conceptual foundation of the empirical study) underpinning our analysis of municipal solid waste (MSW) recycling systems. We begin by describing the recycling rate as the dependent variable and the policy as the treatment variable, followed by a discussion of municipal characteristics as control variables (Section 4.1). Next, we introduce our theoretical framework for the empirical MSW recycling case study, which is based on established waste management theory to capture the interaction between recycling policies, municipal characteristics, and recycling rates (Section 4.2). Finally, we analyze the correlations between these variables to highlight the endogeneity challenges in both policy decisions and uncertain recycling rate (Section 4.3). Together, the data and theoretical framework provide the foundation for the optimization model introduced in Section 5, particularly on how endogenous recycling policies and recycling rates are modeled.

4.1. Data

In this section, we describe our cross-sectional data, which includes the recycling rate as the dependent variable, policy as the treatment variable, and municipal characteristics as control variables:

Recycling Rate (Dependent Variable) and Policy (Treatment Variable) Data. We consider data on recycling policies and recycling rates for 351 Massachusetts municipalities (MassDEP 2022). Some municipalities have multiple or partial policy options that require preprocessing to standardize the data. After preprocessing, a total of 113 municipalities are classified as having curbside recycling services, a total of 102 municipalities are implementing drop-off policies, and a total of 136 municipalities do not have formal recycling programs. In municipalities without a policy, recycling is not assumed to occur.

Figure 2 Recycling Ratio: a) All Municipalities, b) Municipalities by Policy, c) Q-Q plots by Policy

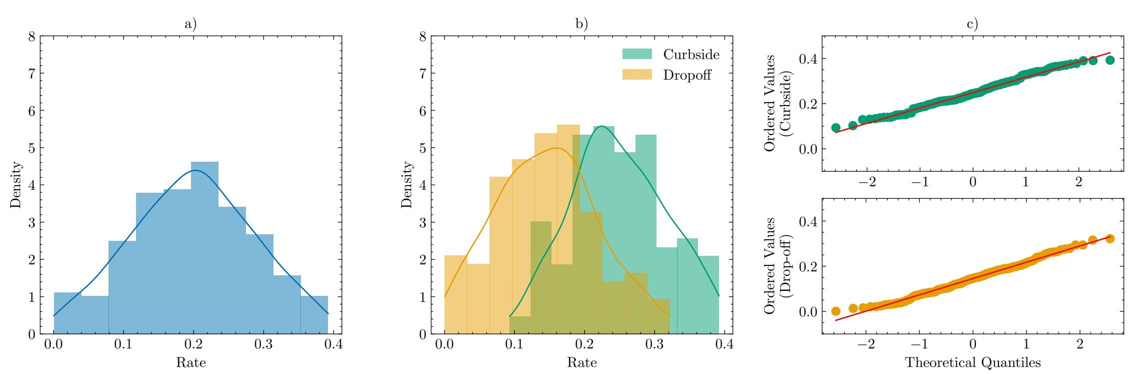

We performed statistical analyses to quickly observe the differences in recycling rates between these policies and identify their distribution (see Figures 2a-b). First, recycling rates for both policies approximate normal distributions, as indicated by the Shapiro-Wilk tests ($p = 0.14$ for curbside;

$p = 0.31$ for drop-off; $p = 0.12$ for all) and the Q-Q plots in Figure 2c). Second, municipalities with curbside collection achieve substantially higher recycling rates ($M = 0.249$, $SD = 0.068$) compared to those using drop-off centers ($M = 0.145$, $SD = 0.072$). A two-sample t-test confirms that this difference is statistically significant ($p = 3.39 \times 10^{-28}$; $t(277) = 12.36$), with a mean difference of 0.103 and SD of 0.008.

Municipal Characteristics Data (Control Variables). To account for heterogeneity across municipalities, we include a range of control variables derived from municipal-level census data (US Census Bureau 2022b) as shown in Table 1. These variables are chosen because they capture socioeconomic and demographic characteristics that influence both policy adoption and recycling behavior, as described in the theoretical framework in the next section.

Table 1 Municipal Characteristics Variables

Variable	Type	Description
Foreign	Socioculture	Percentage of foreign-born persons, 2018-2022
Language	Socioculture	Percentage of persons (5+) speaking a non-English language at home, 2018-2022
White	Socioculture	Percentage of white population, 2018-2022
Education	Ability	Percentage of persons (25+) with a Bachelor's degree or higher, 2018-2022
Poverty	Ability	Percentage of persons living in poverty, 2018-2022
Income	Ability	Median household income (in 2022 dollars), 2018-2022
Older	Demography	Percentage of persons 65+, 2018-2022
Density	Demography	Population per square mile (persons/sqm) , 2020
Owner	Demography	Percentage of owner-occupied housing units, 2018-2022
FamilySize	Demography	Average number of persons per household, 2018-2022

4.2. Theoretical Framework

Our theoretical framework builds on established waste management utility maximization models, as developed by Kinnaman and Fullerton (2000), Sidique et al. (2010), and Starr and Nicolson (2015). These models conceptualize municipal-level recycling outcomes as aggregating individual household waste management decisions. Households seek to maximize their utility from waste disposal choices, balancing financial costs, convenience, social norms, and awareness of recycling benefits within their budget constraints. To better capture the contextual factors influencing recycling behavior, we extend this framework by incorporating social and cultural variables, following Laidley (2013) and Starr and Nicolson (2015). For each municipality $i \in I$, the recycling rate Y_i is modeled as a function of four key factors:

$$Y_i = f(D_i, S_i, B_i, C_i),$$

where:

- D_i (Policy): The municipal policy reflects the relative accessibility, cost, and effort associated with recycling. This is captured through the variable *Curbside*, which reflects the policies that municipalities choose to provide MSW recycling collection services.

- S_i (Socioculture): Social and cultural norms influencing recycling behavior. Following Starr and Nicolson (2015) and Laidley (2013), we include variables such as *Foreign*, *Language*, and *White* to account for cultural diversity and its impact on recycling participation.
- B_i (Ability): Awareness of recycling's benefits and financial ability to afford desired services. These are represented by socioeconomic variables such as *Education*, *Poverty*, and *Income*, which influence both awareness and affordability.
- C_i (Demography): Demographic characteristics that shape recycling behavior. These include variables such as *FamilySize*, *Older*, and *Density*, which affect recycling logistics and residents' participation in recycling.

We must emphasize that we adapt and validate an established theoretical framework from previous studies, confirming its relevance in our context. Our primary contribution lies in using this theoretical framework to prescribe optimal recycling policy options through a novel data-driven methodology. We also build on the conditional independence assumption used in Laidley (2013), similar to the strict exogeneity assumption for municipal variables outlined in Starr and Nicolson (2015) and Sidique et al. (2010). These assumptions allow us to isolate the effects of municipal characteristics on recycling outcomes while accounting for the interdependencies between policy decisions and local contexts.

4.3. Correlation Analysis and Endogenous Uncertainty and Policy

The Pearson correlation matrix (Figure 3) highlights key relationships among the recycling rate, Curbside variable (curbside policy = 1; drop-off policy = 0), and municipal characteristics:

- *Endogenous Policy*: Curbside is strongly correlated with the recycling rate ($r = 0.60$), indicating its pivotal role in influencing performance. However, Curbside also correlates with Language ($r = 0.48$), Foreign ($r = 0.44$), and Density ($r = 0.41$), suggesting potential endogeneity where local characteristics influence policy choice by municipalities.
- *Endogenous Uncertainty*: The positive correlation between recycling rate and Curbside suggests that policy decisions directly impact uncertainty in outcomes. This makes recycling rate uncertainty endogenous for the optimization of policy and RLND.
- *Multicollinear Covariate Relationships*: Income is positively correlated with Education ($r = 0.78$) but negatively correlated with Poverty ($r = -0.61$). Language correlates strongly with Foreign ($r = 0.93$) but inversely with White ($r = -0.92$), highlighting multicollinear demographic patterns within municipalities. This underscores the need for covariate selection.

These findings emphasize the need to account for these two types of endogeneity when modeling the impact of policies on recycling rates.

Figure 3 **Pearson Correlation Matrix of Massachusetts Municipalities Data**

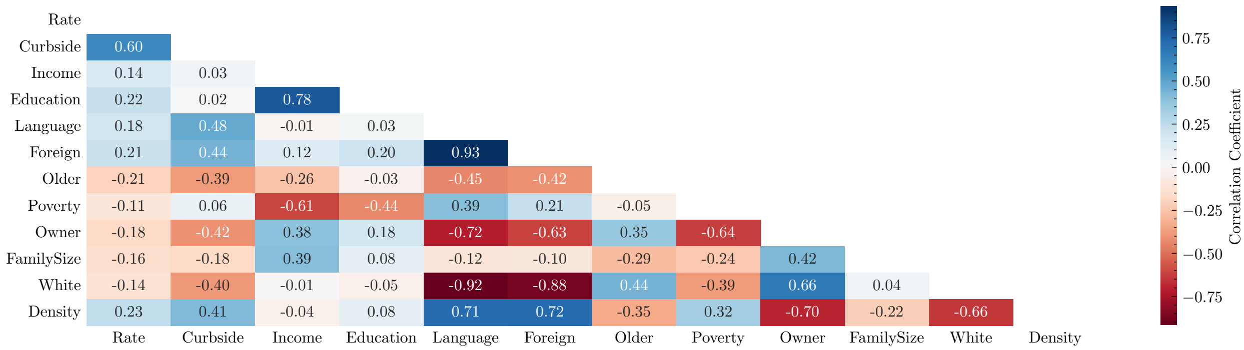

5. Endogenous Contextual Stochastic Optimization Model

This section presents an integrated optimization framework to address the challenges of municipal solid waste (MSW) recycling policy design and reverse logistics network optimization. Our framework is composed of three interconnected components summarized below. First, in Section 5.1, we model the recycling rate as a stochastic outcome influenced by both endogenous policy choices (curbside or drop-off) and municipal characteristics. These relationships are based on the data and theory analyzed in Section 4 and are estimated using the subset of municipalities that implement a recycling policy. In this section, we formally model these relationships as inputs in the optimization model.

Second, in Section 5.2, we provide an expression of collection routing costs using continuum approximation (CA). Here, we extend previous work by linking scalable routing costs estimates to endogenous stochastic recycling rates, municipal waste volumes, vehicle capacities, and geographic characteristics.

Third, in Section 5.3, we formulate a two-stage stochastic RLND problem that integrates the formulations from the previous steps. The first stage determines optimal recycling policies (curbside or drop-off) and activates specific MRFs to serve municipalities. The second stage optimizes material flows between municipalities and MRFs, subject to profitability and capacity constraints. This integration ensures that policy decisions directly influence facility utilization and routing costs.

5.1. Endogenous Recycling Policy and Uncertain Recycling Rate

We consider a set of municipalities I that implement a drop-off or curbside policy. Our goal is to predict the uncertain recycling rate Y_i for each $i \in I$ based on the binary policy assignment D_i and the covariates X_i, which represent municipal characteristics (e.g., income level or education) and are assumed to be independent of the policy assignment. Each municipality is assigned a binary policy $D_i \in \{0,1\}$, where $D_i = 1$ denotes a curbside system and $D_i = 0$ a drop-off system. The parameter τ captures the impact of the curbside system on recycling rates relative to the drop-off system, β captures the covariates' impact, and ϵ_i is zero conditional mean error. Each municipality that implements a recycling policy has two correctly specified potential outcomes:

$$Y_i^0(X_i) = \alpha + X_i^\top \beta + \epsilon_i \quad \text{and} \quad Y_i^1(X_i) = \alpha + X_i^\top \beta + \tau + \epsilon_i,$$

so that the consistent observed outcome

$$\begin{aligned} Y_i(X_i, D_i) &= Y_i^0(X_i)\,(1 - D_i) \; + \; Y_i^1(X_i)\,D_i \\ &= \alpha + X_i^\intercal \beta + \tau D_i + \epsilon_i \text{ where } \mathbb{E}[\epsilon_i | X_i, D_i] = 0. \end{aligned}$$

Here, policy assignment D_i arises through a logistic function of X_i, i.e., $\Pr(D_i = 1 \mid X_i) = G(\zeta X_i + \upsilon_i)$, where $G(b) = (\exp(b))/(1 + \exp(b))$ and $\mathbb{E}[\upsilon_i \mid X_i] = 0$. Here, we observe that D_i is endogenous since it is impacted by exogenous covariates X_i by ζ. Error υ_i is also zero conditional mean.

We impose three sets of assumptions in the following:

Identification and Estimation. This set of assumptions, commonly assumed in causal inference under conditional independence (Angrist and Pischke 2009), allows us to accurately estimate the impact of recycling policies on recycling rates. They ensure that we can appropriately predict Y_i by establishing a valid causal relationship between the policy D_i and the recycling rate. These assumptions help us estimate the policy effect without bias.

1. *Conditional Independence*: $(Y_i^0, Y_i^1) \perp\!\!\!\perp D_i \mid X_i$. Controlling for X_i makes the policy assignment D_i independent of potential outcomes. Therefore, the differences in the recycling rates are solely due to the policy.
2. *Overlap*: $0 < \Pr(D_i = 1 \mid X_i) < 1$ for all X_i. This ensures that every set of characteristics X_i has a non-zero probability of being assigned either curbside ($D_i = 1$) or drop-off ($D_i = 0$), allowing valid estimation of policy effects without exclusive assignment to one policy.
3. *Stable Unit Treatment Value Assumption (SUTVA)*: Each municipality's recycling rate depends solely on its assigned policy, with no spillover effects between municipalities, and the policy is implemented uniformly.

Sparsity. There exists an integer s such that $s\log(p)/n \to 0$ as $n \to \infty$, and $\|\beta\|_0 \leqslant s$, $\|\zeta\|_0 \leqslant s$,where p denotes the number of covariates in X_i. That is, the models depend on at most s covariates, so that $s\log(p)/n \to 0$ as sample size n increases (Belloni et al. 2014).

Modularity. This assumption, common in causal analysis (Neal 2020), requires that (i) the outcome model $Y_i = \alpha + X_i^\top \beta + \tau D_i + \epsilon_i$ remains structurally invariant to changes in the policy D_i, and (ii) the distributions of the covariates X_i and errors ϵ_i are unchanged when D_i is modified.

Additionally, we assume the support of X, Y, D, ϵ are $X \in \mathcal{X} \subseteq \mathbb{R}^{|I| \times |S|}$, $Y \in \mathcal{Y} \subseteq \mathbb{R}^{|I|}$, $D \in \mathcal{D} \subseteq \{0, 1\}^{|I|}$, $\epsilon \in \Xi \subseteq \mathbb{R}^{|I|}$. We further assume that $\mathcal{Y}$ is nonempty closed and convex, which is required for the orthogonal projection onto $\mathcal{Y}$ to be unique and Lipschitz continuous, and $\mathcal{D}$ is nonempty and compact, since it is a finite set. We note that this modeling framework is commonly found in causal inference studies under conditional independence (Angrist and Pischke 2009), and we aim to bridge those studies to prescriptive analytics via ER-SAA (Kannan et al. 2025). We finally note that the recycling policy D_i is endogenous since it is also impacted by covariates X_i. We can also safely assume that municipalities that do not have any policy do not have any recycling volume. Overall, these sets of assumptions allow us to estimate the impact of recycling policies on recycling rates and ensure that this relationship can be reliably incorporated into an optimization model. The policy and rate are incorporated to the contextual model in the next section.

5.2. Contextual Routing Cost Approximation

We extend traditional continuum approximation (CA) methods to account for the variability in recycling rates, which depend on municipal characteristics and policy decisions. Standard CA assumes fixed inputs for routing costs, such as static waste volumes, but this assumption does not hold in our setting. Recycling rates are stochastic outcomes influenced by covariates X_i and binary policy decisions D_i. To address this, we incorporate the stochastic recycling rates $Y_i(x_i, D_i)$, derived from Section 5.1, into the routing cost formulas. This contextual extension ensures that routing costs reflect the dynamic nature of recycling outcomes. The collection routing costs $K^c_{ij}(Y_i(x_i,1))$ for curbside and $K^d_{ij}(Y_i(x_i,0))$ for drop-off, for every municipality i and MRF j, are contextual extensions of the routing cost approximation developed by Winkenbach et al. (2016):

$$K^d_{ij}(Y_i(x_i,0)) = \left\lceil \frac{v^w_i Y_i(x_i,0)}{m^\nu \rho} \right\rceil \left(2c^\delta \delta^L_{ij}\rho + c^\theta \theta^L_{ij}\rho + \frac{c^\nu}{\eta} \right), \tag{1a}$$

$$K^c_{ij}(Y_i(x_i,1)) = c^\delta \nu_{ij}\phi_{ij}\rho\eta \left(2\delta^L_{ij} + \lambda_{ij}\delta^R_i\right) \tag{1b}$$

$$+ c^\theta \nu_{ij}\phi_{ij}\rho\eta \left(\theta^L_{ij} + \lambda_{ij}\theta^R_i\right) \tag{1c}$$

$$+ c^\nu \nu_{ij}, \tag{1d}$$

where the parameters are given by

$$\delta^R_i = \kappa\psi_i\sqrt{\frac{\Omega_i}{\chi_i}}, \tag{1e}$$

$$\pi_i = \frac{v^w_i Y_i(x_i,1)}{\chi_i \rho} \tag{1f}$$

$$\psi_i = \frac{m^\nu}{\pi_i}, \tag{1g}$$

$$\theta^L_{ij} = \theta^S + \frac{2\psi^L_{ij}\delta_{ij}}{\sigma^L}, \tag{1h}$$

$$\theta^R_i = \frac{\delta^R_i}{\sigma^R} + \theta^C, \tag{1i}$$

$$\lambda_{ij} = \begin{cases} \psi_i, & \theta^M \geqslant \theta^L_{ij} + \psi_i \theta^R_i, \\ \dfrac{\theta^M - \theta^L_{ij}}{\theta^R_i}, & \text{otherwise,} \end{cases} \tag{1j}$$

$$\phi_{ij} = \begin{cases} \dfrac{\theta^M}{\theta^L_{ij} + \psi_i \theta^R_i}, & \theta^M \geqslant \theta^L_{ij} + \psi_i \theta^R_i, \\ 1 & \text{otherwise,} \end{cases} \tag{1k}$$

$$\nu_{ij} = \frac{\chi_i}{\lambda_{ij}\phi_{ij}\eta}. \tag{1l}$$

The drop-off routing cost Eq. (1a) consists of 1) rounded up annualized number of linehaul trips required given drop-off rate $Y_i(x_i, 0)$, MSW volume v^w_i, and annualized pickup frequency ρ; 2) total distance-based cost given cost per distance c^δ, line haul distance δ^L_{ij} between municipality i and MRF j, and frequency ρ; 3) total time-based given cost per hour c^θ, line haul duration θ^L_{ij}, and frequency ρ; 4) amortized vehicle cost c^v divided by the number of working days η.

The curbside routing cost also consists of the same last three components: 1) total distance-based costs (Eq. (1b)), which now includes the routing distance for the municipality i, δ^R_i; 2) total time-based costs (Eq. (1c)), which now includes routing duration for the municipality i, θ^R_i; 3) vehicle costs (Eq. (1d)). These costs depend on the number of vehicles needed ν_{ij}, the effective number of routes ϕ_{ij}, and the effective number of stops λ_{ij}.

The rest of the parameters are as follows. Eqs. (1e) are the average inter-stop routing distances for every municipality i, δ^R_i, which consist of distance proportionality factor κ, average circuitry factor ϕ_i, and area Ω_i, and number of households χ_i. Eqs. (1f) calculate the average drop size π_i, dividing MSW volume v^w_i and curbside rate $Y_i(x_i, 1)$ with number of households χ_i and frequency ρ. Eqs. (1g) calculate the maximum number of stops ψ_i that would fit on the vehicle given vehicle capacity m^ν. Eqs. (1h) calculate the line-haul duration θ^L_{ij} which consists of the setup time θ^S and the actual duration of the line-haul trip given circuity ψ^L_{ij}, distance δ_{ij}, and line-haul speed σ^L. Eqs. (1i) calculate the routing duration by dividing average inter-stop routing distance θ^R_i with routing speed σ^R and adding average collection time at every household θ^C. Eqs. (1j) calculate the effective number of stops λ_{ij}, which will be equal to the maximum number of stops ψ_i if total routing time with ψ_i stops and linehaul time does not exceed maximum time θ^M. Otherwise, λ_{ij} will calculate the effective number of stops using the times. Using the same logic, Eqs. (1k) calculate the effective number of routes ϕ_{ij}, which can be greater than one if time allows. Otherwise, it will be one. Finally, Eqs. (1k) calculate the number of vehicles needed ν_{ij} given the number of households χ_i, effective number of stops λ_{ij}, effective number of routes ϕ_{ij}, and number of working days η. Overall, CA has been established as an accurate method to approximate the routing. For numerical studies showing the accuracy of CA, we refer to Janjevic et al. (2019) and Winkenbach et al. (2016).

5.3. Contextual Reverse Logistics Network Design

Given our observed random outcome $Y(X,D)$ in Section 5.1, we can now model the two-stage CSO model for RLND, wherein the first stage, we observe covariates $X = x$ and make policy decisions $D = d$. Specifically, we have three decision variables: $d, g, w \in \mathcal{F}$, where $\mathcal{F}$ is the feasible region. First, for every municipality $i \in I$, d_i is the curbside policy decision, where $d_i = 1$ for curbside and $d_i = 0$ for dropoff. Second, g_i is the decision about whether a municipality i implements a recycling policy, where $g_i = 1$ if it has a recycling policy and $g_i = 0$ if not. Third, for every MRF $j \in J$, $w_j = 1$ if MRF is activated and 0 otherwise. The feasible region $\mathcal{F}$ can then be expressed as:

$$\mathcal{F} = \Big\{(d,g,w) \Big| \, d_i \leqslant g_i \leqslant 1, \, d_i = 1 \, \forall i \in I^c, \, d_i, g_i, w_j \in \{0,1\}, \, \forall i \in I, \, \forall j \in J \Big\}.$$

This region ensures three constraints are fulfilled. First, municipalities can decide on curbside or drop-off with d_i if g_i is activated. Second, municipalities I^c are dedicated as curbside to ensure a minimum recycling rate. Third, all first-stage decision variables are binary.

With these three decision variables, our objective is to solve the approximate solution

$$A^*(x) := \min_{d,g,w \in \mathcal{F}} \quad RLND(d,g,w,Y(x,d)), \tag{2}$$

where the first-stage objective function

$$\begin{aligned} RLND(d,g,w,Y(x,d)) &:= \mathbb{E}_Y[H(d,g,w,Y(x,d)) | X = x, D = d] && \text{(3a)} \\ &= \sum_i b_i^c d_i + b_i^d (g_i - d_i) && \text{(3b)} \\ &+ \sum_i c^e v_i^w t + \mathbb{E}_Y[Q(d,g,w,Y(x,d)) | X = x, D = d]. && \text{(3c)} \end{aligned}$$

The problem setting described in Section 2.2 highlights the need for a comprehensive framework to optimize municipal recycling systems. The key challenge lies in balancing policy decisions, infrastructure investments, and environmental costs while accounting for the stochastic nature of recycling rates and operational complexities. The objective function (2) addresses these challenges by minimizing the total government cost, which includes both the infrastructure costs to implement recycling policies and the external costs associated with the impact of MSW recyclables. The constraints ensure feasibility by linking policy decisions, facility activation, and material flows to operational and environmental requirements.

The objective function (2) minimizes the expected value of the total government cost H given the observation $X = x$ and the decision $D = d$ (Eq. (3a)) consisting of two components. The first is the total costs of the government infrastructure for all municipalities $i \in I$ (Eq. (3b)), where b_i^c and b_i^d are infrastructure costs to invest in curbside and drop-off recycling, respectively. The second is the total external costs (Eq. (3c)), where c^e is the external or environmental cost, v_i^w is the waste volume for every $i \in I$, and t is the recycling target. The first term of Eq. (3c) is constant as it is the total external costs of the waste generated by all municipalities. The second term is the expectation of the second stage model

$$Q(d,g,w,Y(x,d)) := \min_z \; -c^e\left(\sum_{i\in I}\sum_{j\in J}(z_{ij}^c + z_{ij}^d)\right) \tag{4a}$$

$$\text{s.t.} \sum_{i\in I} U_{ij}^c(Y_i(x_i,1))z_{ij}^c + U_{ij}^d(Y_i(x_i,0))z_{ij}^d - c_j w_j \geqslant 0, \qquad \forall j \in J, \tag{4b}$$

$$\sum_{i\in I}\sum_{i\in J} U_{ij}^c(Y_i(x_i,1))z_{ij}^c + U_{ij}^d(Y_i(x_i,0))z_{ij}^d - c_j w_j \geqslant \bar{P}, \tag{4c}$$

$$\sum_{j\in J} z_{ij}^c \leqslant M d_i, \qquad \forall i \in I, \tag{4d}$$

$$\sum_{j\in J} z_{ij}^d \leqslant M(g_i - d_i), \qquad \forall i \in I, \tag{4e}$$

$$\sum_{j\in J}(z_{ij}^c + z_{ij}^d) = v_i^w Y_i(x_i,d_i), \qquad \forall i \in I, \tag{4f}$$

$$\sum_{i\in I}(z_{ij}^c + z_{ij}^d) \leqslant m_j^f w^j, \qquad \forall j \in J, \tag{4g}$$

$$z_{ij}^c, z_{ij}^d \geqslant 0, \qquad \forall i \in I, \forall j \in J, \tag{4h}$$

where unit profit for curbside and drop-off equations

$$U_{ij}^c(Y_i(x_i,1)) = r^s p^r + r_i^t - c^p - (1-p^r)c^l - \frac{K_{ij}^c(Y_i(x_i,1))}{v_i^w Y_i(x_i,1)}, \tag{4i}$$

$$U_{ij}^d(Y_i(x_i,0)) = r^s p^r + r_i^t - c^p - (1-p^r)c^l - \frac{K_{ij}^d(Y_i(x_i,0))}{v_i^w Y_i(x_i,0)}. \tag{4j}$$

The objective function (4a) maximizes the total environment cost savings from curbside and dropoff recycling flows z_{ij}^c and z_{ij}^d of every county i to MRF j, where c^e is the environmental cost per ton. Constraints (4b) ensure every MRF j has a nonnegative profit. The profit consists of the total variable profit from multiplying the flows with the unit profit U_{ij}^c (curbside) and U_{ij}^d (dropoff) and subtracting it with MRF fixed cost c_j^f. Constraint (4c) ensures that the total profit for all MRFs are greater than the minimum total profit $\bar{P}$. Using large constant M, constraints (4d) ensure curbside flows are greater than zero only when the curbside policy is activated ($d_i = 1$) for every municipality i. The inverse is true for the drop-off policy in Constraints (4e) ($d_i = 0$). Notice that both curbside and drop-off policies require policy activation $g_i = 1$. Otherwise, no recycling policy is implemented in the municipality i. Constraints (4f) connects the endogenous random outcome $Y_i(x_i, d_i)$ and the recycling flow for every municipality i, v_i^w with the total flow of curbside and drop-off. Together with Constraints (4d) and (4e), we observe that when the curbside policy is activated ($d_i = 1$), only curbside recycling flows z_{ij}^c will be activated. The same goes for the drop-off policy ($d_i = 0$). Constraints (4g) ensures capacity of every MRF j given capacity m_j^f and links flows with MRF activation. Constraints (4h) ensures nonnegative flows. The unit profit Equations (4i) and (4j) consists of scrap sales revenue r^s given processing rate p^r, tipping collection revenue r_i^t, processing cost c^p, and landfill cost c^l and the collection routing costs $K_{ij}^c(Y_i(x_i, 1))$ (for curbside) and $K_{ij}^d(Y_i(x_i, 0))$ (for drop-off), which we unitize by their volumes and describe in Section 5.2.

6. Solution Approach

To solve the problem $A^*(x)$ in Eq. (2), we require knowledge of the distributions governing the recycling rate. In practice, however, we have access only to historical data on Massachusetts's municipal recycling supply, policies, and characteristics (MassDEP 2019, 2022, US Census Bureau 2022a). To bridge this gap, we adopt the ER-SAA approach (Kannan et al. 2025, Zhu et al. 2024) combined with PDS RLasso methods (Belloni et al. 2012, 2014). This framework allows us to predict recycling supply and subsequently optimize the RLND problem using observed data.

6.1. Empirical Residuals-based Sample Average Approximation Approach

The ER-SAA approach (Kannan et al. 2025) involves several steps. First, we train a machine learning model on historical data to predict the recycling rate (Y_i). Second, we compute residuals (the differences between actual values and predictions) to capture the uncertainty not explained by the model. Finally, we use both predictions and empirical residuals to construct scenarios for the stochastic programming model and solve it using SAA method. However, as we later observe, this approach alone is insufficient due to the endogeneity of policies and recycling rate uncertainties.

We can formalize this approach as follows. Let data $\Delta_n := \{(x^l, d^l, y^l)\}_{l=1}^n$ denote joint observations of municipal characteristics X, recycling policy D, and recycling supply Y for n municipalities, where for each municipality data observation $l \in L$, x^l represents demographic and socioeconomic characteristics, d^l indicates the historical collection policy (with curbside represented by 1 and dropoff by 0), and y^l represents the observed recycling supply. These observations are obtained from Massachusetts municipal recycling data and census information. We previously model that the true relationship between recycling supply Y, municipality characteristics X, and collection policies D is given by

$$Y = f^*(X, D) + \epsilon,$$

where $f^*(X, D)$ represents the true conditional mean recycling supply given characteristics X and policies D, and ϵ is a zero-mean regression error independent of both X and D, as described in Section 5.1. In the full-information case where f^* is known, given Δ_n, the Full-Information Sample Average Approximation (FI-SAA) for RLND is formulated as:

$$A_n^{FI}(x) := \min_{d,g,w \in \mathcal{F}} RLND_n^{FI}(d, g, w, f^*(x, d) + \epsilon)$$

$$= \min_{d,g,w \in \mathcal{F}} \frac{1}{n} \sum_{l=1}^{n} H(d,g,w,f^*(x^l,d^l) + \epsilon^l),$$

where $\epsilon^l = y^l - f^*(x^l, z^l)$. Since f^* is unknown in practice, we estimate it via regression methods on Δ_n, yielding the estimator $\hat{f}_n$. This leads to the ER-SAA formulation for RLND:

$$\begin{aligned} A_n^{ER}(x) &:= \min_{d,g,w \in \mathcal{F}} RLND_n^{ER}(d,g,w,\mathrm{proj}_{\mathcal{Y}}(\hat{f}_n(x,d) + \hat{\epsilon}_n)) \\ &= \min_{d,g,w \in \mathcal{F}} \frac{1}{n} \sum_{l=1}^{n} H\Big(d,g,w,\mathrm{proj}_{\mathcal{Y}}\Big(\hat{f}_n(x^l,d^l) + \hat{\epsilon}_n^l\Big)\Big), \end{aligned}$$

where the empirical residuals are computed as $\hat{\epsilon}_n^l = y^l - \hat{f}_n(x^l, d^l)$, and $\mathrm{proj}_{\mathcal{Y}}(y)$ represents the orthogonal projection onto the support of feasible recycling rate $\mathcal{Y}$. This projection ensures that our predictions remain feasible (e.g., nonnegative), thereby preventing potential infeasibility issues in the second-stage problem (Zhu et al. 2024, Kannan et al. 2025). However, this dependence of $\hat{f}_n(x,d)$ on the decision variable d renders the ER-SAA formulation intractable (Zhu et al. 2024), motivating the need for a method that ensures $\hat{f}_n$ depends only on the exogenous covariates x.

6.2. Post Double Selection with Rigorous Lasso

As discussed in the previous section, the endogenous uncertainty in our model presents a challenge for standard ER-SAA by (Kannan et al. 2025). Currently, standard ER-SAA approaches assume decision-independent uncertainties. However, our review of the literature (e.g., Kinnaman and Fullerton (2000)) and the correlation analyses in Section 4.3 reveal that recycling rates depend on both policy decisions and municipal characteristics. While Zhu et al. (2024) addresses this using distributionally robust optimization, we instead modify the ER-SAA method itself.

In the prediction step of ER-SAA, we choose the PDS method developed by Belloni et al. (2014) for two reasons. First, it addresses multicollinear covariates identified in Section 4.3 and selects covariates sparsely. Second, it effectively handles endogenous policy and uncertainty, placing the true regression function $f^*(x,d)$ and estimating it consistently using $\hat{f}_n(x,d)$. If we use a regression model with only the policy variable (d), it fails to properly address the endogeneity of policy because the policy variable itself is correlated with other municipal characteristics. Conversely, employing a full OLS regression model with all available covariates fails to adequately manage the multicollinearity among highly correlated predictors. We further support these observations with numerical evidence presented in Section 7.1.

PDS executes two Lasso regressions at a high level: one focusing on policy decisions and the other on recycling rates. This dual regression strategy selects the relevant control variables and predictors. The method then combines the selected variables from both regressions in a final OLS regression, thereby mitigating the bias of the omitted variables and enabling valid inference about the effects of the policies. To further enhance the Lasso regression steps, we incorporate Rigorous Lasso (Belloni et al. 2012). This refinement employs theoretically grounded penalty levels, ensuring the selection of genuinely relevant variables and improving estimate consistency in a sparse manner.

6.3. Our Combined Approach

Our final solution strategy combines ER-SAA with PDS RLasso. It consists of two main steps that address endogeneity in recycling policy and uncertainty while estimating the conditional mean function $f^*(x,d)$:

Algorithm 1: ER-SAA with PDS RLasso

Step 1. Post Double Selection with Rigorous Lasso:

- **a.** *Treatment Covariate Selection:* Apply RLasso to regress the treatment variable d on the covariates x. This step calculates the coefficients $\hat{\zeta}^{L_1}$ and the set of non-zero covariates $\hat{S}_D := \{j \in \{1,\dots,|S|\} : \hat{\zeta}_j^{L_1} \neq 0\}$.
- **b.** *Outcome Covariate Selection:* Similarly, perform RLasso to regress the outcome variable Y_l on the covariates X_l, yielding non-zero coefficients $\hat{\beta}^{L_1}$ and the associated set of covariates $\hat{S}_Y := \{j \in \{1,\dots,|S|\} : \hat{\beta}_j^{L_1} \neq 0\}$.
- **c.** *Union of Selected Covariates:* Form the union of the selected covariates, $\hat{S} := \hat{S}_D \cup \hat{S}_Y$.

Step 2. Empirical Residuals-based Sample Average Approximation:

- **a.** *Conditional Mean Estimation:* Estimate the function $f^*(x,d)$ using the selected covariates x_s corresponding to $\hat{S}$. Specifically, we fit the model
$$\hat{f}_n(x,d) = \hat{\alpha} + x_s^{\intercal}\hat{\beta} + \hat{\tau} d.$$
- **b.** *Residual Computation:* For each sample $l \in [n]$, compute the residuals
$$\hat{\epsilon}_n^l = y^l - \hat{f}_n(x^l, d^l).$$
- **c.** *Proxy Sample Generation and Optimization:* Use the projected estimates $\{\mathrm{proj}_{\mathcal{Y}}(\hat{f}_n(x^l, d_i) + \hat{\epsilon}_n^l)\}_{l=1}^n$ given d_i as policy for municipality $i \in I$ as proxy samples for recycling supply in the optimization problem. The empirical ER-SAA formulation then becomes:
$$A_n^{ER}(x) = \min_{d,g,w \in \mathcal{F}} \frac{1}{n} \sum_{l=1}^{n} H\Big(d, g, w, \mathrm{proj}_{\mathcal{Y}}\big(\hat{f}_n(x^l, d^l) + \hat{\epsilon}_n^l\big)\Big).$$

With this approach, we can effectively replace $Y_i(x_i, d_i)$ in Eqs. (2), (3a)-(3c), (4a)-(4h), with $\mathrm{proj}_{\mathcal{Y}}\big(\hat{f}_n(x_i^l, d_i) + \hat{\epsilon}_n^l\big) = \mathrm{proj}_{\mathcal{Y}}\big(\hat{\alpha} + x_i^{l\intercal}\hat{\beta} + \hat{\tau} + \hat{\epsilon}_n^l\big) d_i + \mathrm{proj}_{\mathcal{Y}}\big(\hat{\alpha} + x_i^{l\intercal}\hat{\beta} + \hat{\epsilon}_n^l\big)(1 - d_i)$ for every municipality $i \in I$ and observation sample $l \in L$. The following proposition establishes that, under our modeling assumptions, we can correctly identify the true conditional mean function and that our estimator converges to the true parameters as the sample size increases.

Proposition 1 (Endogenous to Exogenous)**.** *Under assumptions in Section 5.1:*

1. *The true conditional mean is identifiable as*
$$f^*(X_i, D_i) = \mathbb{E}[Y_i \mid X_i, D_i] = \alpha + X_i^{\intercal}\beta + \tau D_i.$$
2. *The estimator $\hat{f}_n(x,d)$ is consistent; that is, as $n \to \infty$, we have $\hat{\alpha} \xrightarrow{p} \alpha$, $\hat{\beta} \xrightarrow{p} \beta$, and $\hat{\tau} \xrightarrow{p} \tau$.*

Remark 1. *Given Proposition 1, the asymptotic optimality and finite-sample convergence results established for exogenous ER-SAA in Kannan et al. (2025) extend naturally to our setting with endogenous uncertainties when using ER-SAA combined with PDS RLasso.*

We effectively mitigate the bias introduced by endogeneity in the recycling policy by appropriately selecting relevant covariates through PDS RLasso and assumptions in Section 5.1. This adjustment allows our data-driven estimator to behave as if the uncertainties were exogenous. As a result, the asymptotic optimality and finite-sample convergence properties that apply in an exogenous setting naturally extend to our endogenous framework.

7. Massachusetts MSW Recycling Case Study

This section presents the results of our analysis of the Massachusetts municipal solid waste (MSW) recycling system. We begin by evaluating the performance of regression models in Section 7.1, followed by an assessment of CSO results in Section 7.2. Key insights include the value of incorporating covariate data and the policy adjustments achieved through CSO. In the E-companion (EC), we additionally discuss data on the economic structure of MRFs and public authorities in Section EC.2 and the current recycling network infrastructure and operations data in Section EC.3. Finally, we describe how these municipal characteristics differ across municipalities by the policies prescribed by the CSO solution in Section EC.4.

7.1. Regress: Predicting Recycling Rate

As discussed in Section 6.2, we select the PDS RLasso approach primarily because it effectively addresses policy endogeneity and mitigates multicollinearity among correlated covariates. We numerically compare three regression models: a baseline model with only the policy variable (Curbside), a full OLS model with all available covariates, and our PDS RLasso approach. Table 2 presents the comparative results which we discuss in the next Section.

7.1.1. OLS Regression Models Results First, we observe that the PDS RLasso model achieves the best adjusted $R^2 = 0.410$, confirming its effectiveness in capturing key predictors while addressing multicollinearity. For robustness, we also show in Section EC.5 in the E-companion that our prediction performance is also superior to other machine learning models. As we optimize for the optimal recycling policy, the public authorities can only change the policy for municipalities and not their characteristics. Finally, in the later optimization results, we use the OLS model with only the policy variable for CSO solution without covariate data and PDS RLasso for CSO solution with covariate data. Second, we further observe that the significant predictors include Curbside, Education, and Language across the models:

Table 2 Ordinary Least Squares (OLS) Regression Estimates for Massachusetts Recycling Rate

	OLS with Curbside Only	OLS with All Variables	OLS with PDS RLasso
Intercept	0.145***	0.215†	0.144**
	(0.006)	(0.124)	(0.047)
Curbside	0.103***	0.109***	0.112***
	(0.008)	(0.011)	(0.010)
Income		-3.184e-07	
		(2.51e-07)	
Education		0.148**	0.112***
		(0.046)	(0.026)
Language		-0.160	-0.175***
		(0.113)	(0.050)
Foreign		0.018	
		(0.124)	
Older		-0.100	-0.010
		(0.107)	(0.076)
Poverty		-0.139	
		(0.156)	
Owner		0.000	-0.044
		(0.076)	(0.052)
FamilySize		-0.019	
		(0.023)	
White		-0.012	
		(0.094)	
Density		1.705e-06	1.909e-06
		(2.25e-06)	(2.13e-06)
R^2	0.357	0.431	0.423
Adjusted R^2	0.355	0.408	0.410
F-statistic	152.0	21.10	33.31
Prob (F-statistic)	4.25e-28	1.61e-30	5.69e-30
Durbin-Watson	1.920	1.962	1.970
Prob (Jarque-Bera)	0.127	0.506	0.553

$^{\dagger}p < .1, ^{*}p < .05, ^{**}p < .01, ^{***}p < .001.$

- The Curbside variable, which represents Curbside policy when compared to Drop-off, shows a consistent and significant positive effect on recycling rates across all three models: 10.3% ($p <$ .001) for Curbside only, 10.9% ($p < .001$) for all variables. The PDS RLasso model estimates a higher Curbside effect at 11.2% points ($p < .001$). This consistency among models suggests that Curbside is a robust predictor of increased recycling rates.
- Education emerges as another significant predictor in both the All Variables and PDS RLasso models. In the All Variables model, 1% increase in Education is associated with a 0.148% increase in recycling rates ($p < .01$). The PDS RLasso model estimates a lower effect, with a 0.112% ($p < .001$). We note that Education is highly positively correlated with Income and highly negatively correlated with Poverty, which could explain why PDS RLasso model

discards these covariates. This suggests that higher levels of education (and by proxy higher income and lower poverty) in a community are strongly associated with higher recycling rates.

- Language variable shows a significant negative effect in the PDS RLasso model, with 1% increase associated with a 0.175 % decrease in recycling rates ($p < .001$). We also note that Language is highly positively correlated with Foreign and highly negatively correlated with White, which could explain why the PDS RLasso model discards these covariates. This suggests that higher rates of foreign language speakers (and, by proxy, higher rates of foreigners and lower white persons) are associated with lower recycling rates.

7.1.2. Prediction and Residuals with and without Covariate Data We observe from Figures 4 and 5 that covariate data captures more heterogeneity of the recycling rate distribution.

- Without covariates (Figures 4a-b), the model produces single point predictions of 0.248 for Curbside and 0.145 for Drop-off (from Table 2). This overly simplistic outcome arises because the model treats every municipality alike, ignoring key demographic or socioeconomic differences. When we add covariate information (Figures 4c-d), the prediction distributions become more dispersed, matching the spread and shape of the ground truth more closely for both Curbside and Drop-off. However, we do not fully capture the complete distribution of observed recycling rates, even with covariates. This gap stems partly from irreducible uncertainties discussed in Section 2.2, i.e. randomness or unmeasured factors that remain unexplainable given the available covariate data. While municipal characteristics capture a substantial portion of the underlying variation, some inherent randomness inevitably persists.
- The residual analysis confirms the superior performance of the covariate model. Without covariates (Figures 5a-b), residuals form two distinct vertical clusters in the fitted values of Drop-off and Curbside, which are the same values from Table 2 and Figures 4a-b. Including covariates (Figures 5c-d) produces more randomly scattered residuals with a more symmetric distribution. Both models pass the Breusch-Pagan test ($p = 0.4303$ and $p = 0.2079$ respectively), confirming homoscedasticity. The Shapiro-Wilk tests ($p = 0.0735$ and $p = 0.2860$) indicate normally distributed residuals in both cases. When comparing 5b) and 5d), the narrower residuals spread in the model with covariates also shows improved accuracy.

7.2. Reverse and Recycle: Contextual Stochastic Optimization Results

This section presents the results of applying contextual stochastic optimization (CSO) to Massachusetts' municipal solid waste recycling system. We first demonstrate the improvements from the CSO solution and show how the incorporation of municipal characteristics into optimization decisions leads to significant improvements in both environmental and economic outcomes.

Figure 4 **Regression Prediction Distribution without Covariate Data (a) Curbside and b) Drop-off Regression Prediction Distribution with Covariate Data (c) Curbside and d) Drop-off**

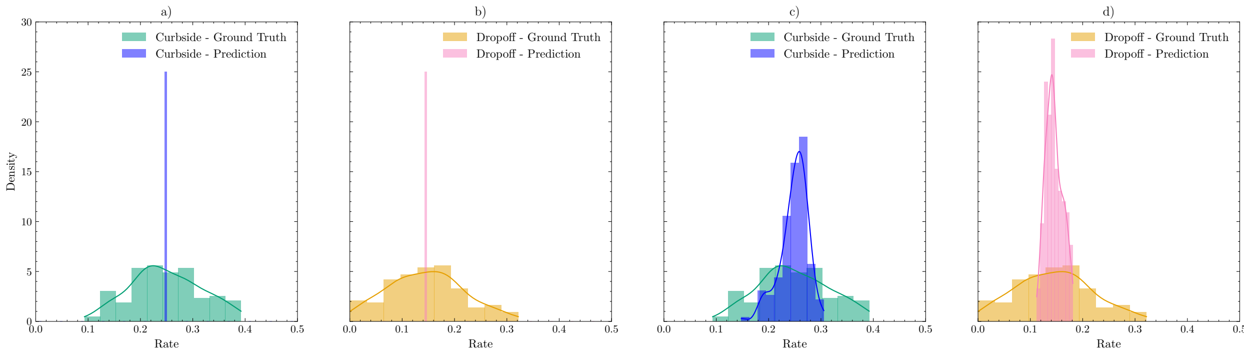

Figure 5 **Regression Residuals without Covariate Data (a) Residuals vs. Fitted Values and b) Histogram Regression Residuals with Covariate Data (c) Residuals vs. Fitted Values and d) Histogram**

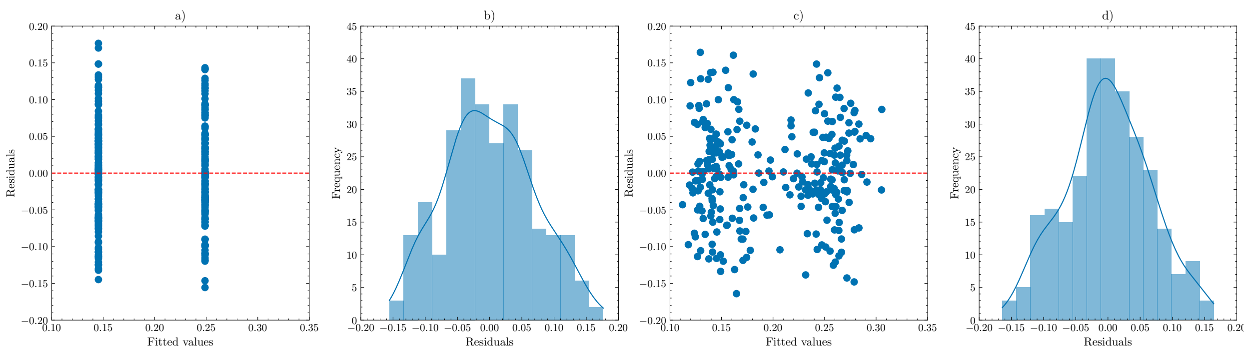

Figure 6 **Current and Candidate MRFs (in Green and Blue respectively) and Municipalities Center Points**

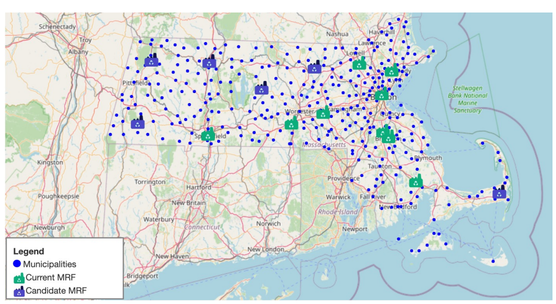

Currently, there are nine operational MRFs in MA, and we consider six additional candidate MRFs in uncovered areas, as shown in Figure 6. To demonstrate the benefit of CSO solution, we compare the following network configurations (see Table 3 and Figures 7, 8, 9, and 10):

- **Current**: The existing network configuration and policy assignments based on Recycle Smart MA (2023) data, with nine operational MRFs.
- **Baseline**: The network is obtained by considering the current policies but re-optimizing the facility location, allocation, and routing decisions.
- **CSS**: The network is the CSO solution with covariate data and PDS RLasso prediction model.

- **CSS-NC**: The network is the CSO solution without using covariate data, whose prediction model is the OLS with only the policy variable.
- **CSS-NC-E**: The evaluation of the CSS-NC solution using uncertainty realizations that incorporate covariate data, which yields results similar network to CSS-NC.

Table 3 Comparison Between Current, Baseline, and CSO Networks

Metric	Current	Baseline	CSS	CSS-NC	CSS-NC-E
Infrastructure Cost ($M)	16.5	16.5	30.9	27.3	27.3
Change from Baseline (%)	-	-	*+87%*	*+65%*	*+65%*
Environmental Cost ($M)	75.3	75.3	40.8	41.9	46.2
Change from Baseline (%)	-	-	*-46%*	*-44%*	*-39%*
Total Government Cost ($M)	91.8	91.8	71.7	69.2	73.5
Change from Baseline (%)	-	-	*-22%*	*-25%*	*-20%*
Total MRF Profit ($)	(5.2)	1.9	7.5	8.3	0.4
Change from Baseline (%)	*-378.4%*	-	*+303%*	*+345%*	*-80%*
Rate (%)	13.7	13.7	21.2	20.9	20.0
Change from Baseline (%)	-	-	*+55%*	*+53%*	*+46%*
MRFs Opened (Out of 15)	9	8	12	12	12
Municipalities w/ Curbside	113	113	148	150	150
Municipalities w/ Dropoff	102	102	171	114	114
Municipalities w/ No Policy	136	136	32	87	87

From the comparison, we describe our findings in the following paragraphs:

Baseline Benchmark: We construct a baseline scenario by re-optimizing the network while keeping recycling policies constant, in order to isolate the effects of facility reallocation from those driven by policy changes. The MRF profitability shifts from a $5.2M loss to a $1.9M profit.

Improvements from CSO: As shown in Table 3, design obtained through CSS and CSS-NC both deliver substantial improvements over the baseline. CSS design increases infrastructure costs by 87% (from $16.5M to $30.9M) but reduces environmental costs by 46% (from $75.3M to $40.8M), resulting in a 22% decrease in total government costs (from $91.8M to $71.7M). Similarly, CSS-NC design increases infrastructure costs by 65% but cuts environmental costs by 44%, resulting in a reduction of 25% in total government costs. Both designs dramatically improve MRF profitability, with CSS increasing profits by 303% (from $1.9M to $7.5M) and CSS-NC by 345% (to $8.3M). Recycling rates improve significantly: CSS achieves a 55% increase (to 21.2%), CSS-NC delivers a 53% increase (to 20.9%).

We can observe the network configurations obtained with these two approaches (Figures 9 and 10). The number of MRFs increases from 8 to 12 with an identical layout, likely due to having only 15 MRFs available. From Figures 11 and 12, we also observe that new MRFs are strategically placed close to previously underserved areas with greater potential to recycle (e.g., in western

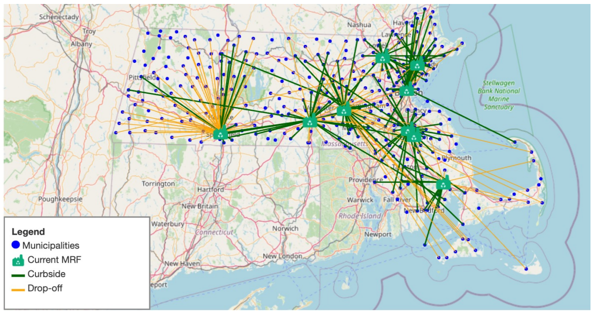

Figure 7 **Current Network**

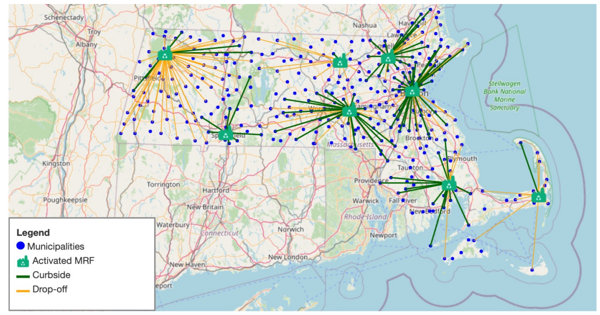

Figure 8 **Baseline Network**

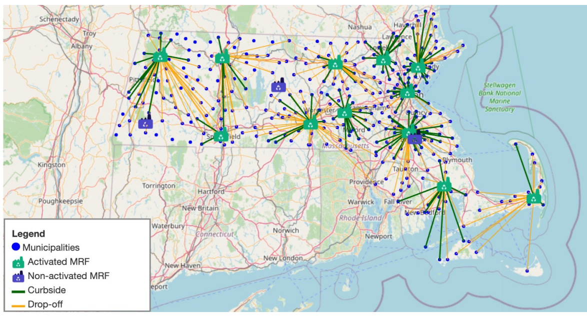

Figure 9 **CSS Policy Network**

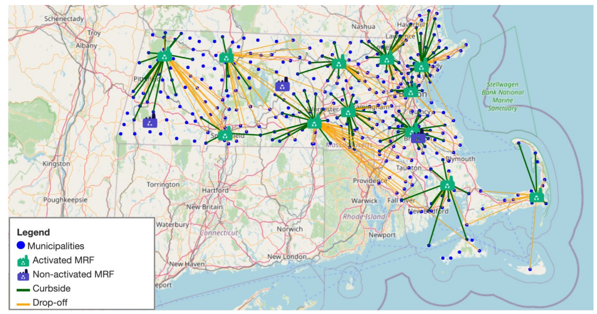

Figure 10 **CSS-NC Policy Network**

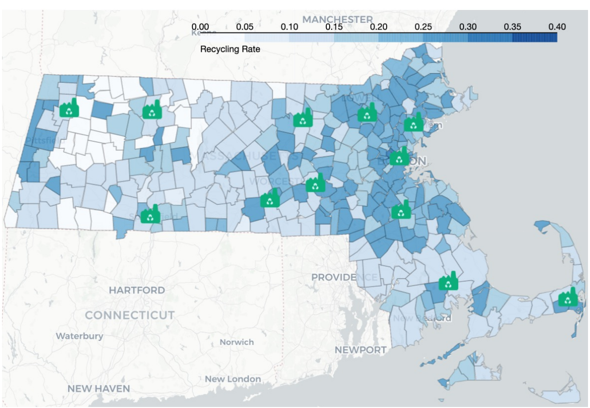

Figure 11 **CSS Recycling Rate**

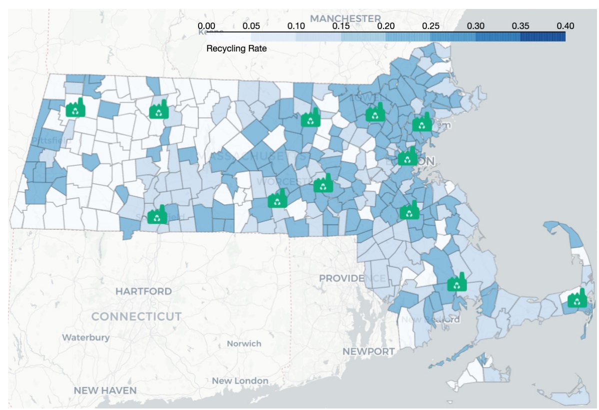

Figure 12 **CSS-NC Recycling Rate**

Massachusetts). The proximity of MRFs to these areas, along with their increased recycling supply, results in profitable MRF operations, illustrating the alignment between policy and network design.

Policy Comparison Table 4 summarizes the policy changes between CSS and CSS-NC. First, we observe that 253 out of 351 municipalities (72%) retain the same policy across both approaches. This alignment is expected since prior policy assignments strongly influence current recommendations. However, several notable differences emerge. The most significant policy shift involves municipalities assigned Drop-off under CSS but receiving different policies in CSS-NC. Specifically, 56 municipalities originally without any policy in CSS-NC were upgraded to Drop-off in CSS, primarily due to their moderate density (255.4 persons/sqm). Similarly, 20 municipalities initially designated for Curbside in CSS-NC were reassigned to Drop-off under CSS, reflecting comparable density levels (277.7 persons/sqm). Another notable adjustment involves municipalities reassigned to Curbside under CSS. Nineteen municipalities initially recommended for Drop-off by CSS-NC shift to Curbside in CSS because of their significantly higher density (3594.6 persons/sqm), making curbside collection economically viable. In summary, CSS-NC employs a uniform recycling rate estimate, which does not fully capture the municipal heterogeneity. In contrast, CSS uses detailed municipal characteristics, enabling more tailored and economically beneficial recycling policy decisions, as detailed in the following paragraph.

Table 4 Policy Comparison Between CSS and CSS-NC Approaches

CSS Policy	CSS-NC Policy	# Municipalities	Density (persons/sqm)	Area (sqm)
Curbside	Curbside	128	2342.8	17.3
	Drop-off	19	3594.6	14.9
	No Policy	1	81.3	5.4
Drop-off	Curbside	20	277.7	29.9
	Drop-off	95	1124.1	23.9
	No Policy	56	255.4	28.3
No Policy	Curbside	2	13.5	17.5
	Drop-off	-	-	-
	No Policy	30	48.1	28.1

Value of Covariate Data: Comparing CSS with CSS-NC-E reveals the importance of incorporating municipal characteristics into decision-making. While CSS-NC appears promising on paper, its actual performance when evaluated with covariate data (CSS-NC-E) shows significant overestimation of benefits in Table 3. Most striking is the MRF profitability: CSS-NC predicts \$8.3M, but CSS-NC-E achieves only \$0.4M—a dramatic 95% reduction. Environmental costs in CSS-NC-E (\$46.2M) are 10% higher than predicted by CSS-NC (\$41.9M), and total government costs are 6% higher (\$73.5M vs. \$69.2M). The recycling rate also falls from the predicted 20.9% to 20.0%. These disparities occur because CSS-NC makes policy recommendations without considering municipal characteristics, leading to mismatches between policies and local conditions. CSS avoids these pitfalls, resulting in more realistic projections and better overall performance.

8. Concluding Discussion

This study addresses the design and operation of the Massachusetts Municipal Solid Waste (MSW) recycling system by Contextual Stochastic Optimization (CSO). By explicitly incorporating the endogeneity of recycling policies and the complexity of routing into a reverse logistics network design (RLND), our work addresses the critical interdependencies between policy, network design and routing. In doing so, we offer a data-driven framework that both explains and prescribes optimal recycling policies.

A notable methodological contribution of our work lies in enabling the Empirical Residuals-based Sample Average Approximation (ER-SAA) framework to handle endogenous policies and uncertainties. In this scenario, the recycling policy itself depends on municipal characteristics. By employing PDS RLasso, we effectively select the relevant covariates for the equations of treatment (that is, curbside policy) and result (that is, recycling rate). Proposition 1 shows that, under reasonable assumptions, this approach transforms the endogenous setting into one where the policy uncertainties can be treated as if they are exogenous. In practice, once the critical covariates have been identified, our ER-SAA formulation can apply standard causal inference technique as prediction and integrate it with optimization to produce policy insights. As a result, the proposed solution captures the nuanced interdependencies between municipal characteristics and policy choices.

Although the notion that leveraging the heterogeneity of municipal characteristics and adjusting policies could naturally lead to increased recycling is intuitive, it is unclear whether this can also be done without increasing cost. The endogeneity between recycling policies and municipal characteristics, integrated with network design and routing decisions, makes it far from obvious how this trade-off will materialize. Our contribution lies in developing a method that transforms these endogenous uncertainties into an exogenous framework, thereby making the optimization problem tractable and the trade-off explicit. We note that even in cases where our approach may not be possible, integrating prediction and optimization can still improve overall performance, as demonstrated by Bertsimas and Kallus (2023) and Fernández-Loría and Provost (2022).

From a predictive or explanatory standpoint, our Post Double Selection with Rigorous Lasso (PDS RLasso) model demonstrates superior performance over OLS or even more complex machine learning approaches, such as the new foundation model TabPFN by Hollmann et al. (2025). By isolating covariates that are relevant for both treatment (curbside vs. drop-off) and outcome (recycling rate), PDS RLasso yields higher adjusted values R^2, mitigates multicollinearity, and provides actionable insights into how municipal characteristics (e.g., education and language nexuses) shape recycling rates. Public authorities can swiftly adjust recycling policies, but factors such as age and education are much harder to change.

From an optimization standpoint, through extensive numerical experiments on Massachusetts' municipal data, we find that our optimal policy and network can increase the overall recycling rate by 7.5%, reduce public authorities' costs by $20.1M, and add $12.7M to the profitability of MRF operators. Incorporating heterogeneity through covariate data into the CSO not only enhances the accuracy but also avoids overstatements in benefits. Notably, while it is clear that adjusting recycling policies and accounting for demographic heterogeneity through municipal characteristics can improve recycling outcomes, it is not immediately evident that doing so will also yield cost savings. Indeed, our findings show that the proposed integrated approach not only increases recycling rates but also reduces overall public costs and increases MRF profits, demonstrating that environmental goals and economic efficiency can go together. Finally, we note the spatial impact of network design, which allows improvement in recycling rate and economic efficiency. Overall, we address the limitations of MSW recycling policy and RLND by developing a new data-driven methodology that can be extended to other stochastic programming problems for further development.

References

Angelelli E, Speranza M (2002) The application of a vehicle routing model to a waste-collection problem: two case studies. *Journal of the Operational Research Society* 53(9):944–952.

Angrist JD, Pischke JS (2009) *Mostly harmless econometrics: An empiricist's companion.*

Anshassi M, Townsend TG (2023) The hidden economic and environmental costs of eliminating kerb-side recycling. *Nature Sustainability* 6(8):919–928.

Ban GY, Rudin C (2019) The big data newsvendor: Practical insights from machine learning. *Operations Research* 67(1):90–108.

Basuhi R, Bhuwalka K, Roth R, Olivetti EA (2024) Evaluating strategies to increase pet bottle recycling in the united states. *Journal of Industrial Ecology* 28(4):916–927.

Basuhi R, Moore E, Gregory J, Kirchain R, Gesing A, Olivetti EA (2021) Environmental and economic implications of us postconsumer plastic waste management. *Resources, Conservation and Recycling* 167:105391.

Belloni A, Chen D, Chernozhukov V, Hansen C (2012) Sparse models and methods for optimal instruments with an application to eminent domain. *Econometrica* 80(6):2369–2429.

Belloni A, Chernozhukov V, Hansen C (2014) Inference on treatment effects after selection among high-dimensional controls. *Review of Economic Studies* 81(2):608–650.

Beltrami EJ, Bodin LD (1974) Networks and vehicle routing for municipal waste collection. *Networks* 4(1):65–94.

Bertsimas D, Kallus N (2020) From predictive to prescriptive analytics. *Management Science* 66(3):1025–1044.

Bertsimas D, Kallus N (2023) The power and limits of predictive approaches to observational data-driven optimization: The case of pricing. *INFORMS Journal on Optimization* 5(1):110–129.

Bing X, Bloemhof JM, Ramos TRP, Barbosa-Povoa AP, Wong CY, van der Vorst JG (2016) Research challenges in municipal solid waste logistics management. *Waste Management* 48:584–592.

Bing X, Bloemhof-Ruwaard JM, Van Der Vorst JG (2014a) Sustainable reverse logistics network design for household plastic waste. *Flexible Services and Manufacturing Journal* 26:119–142.

Bing X, de Keizer M, Bloemhof-Ruwaard JM, van der Vorst JG (2014b) Vehicle routing for the eco-efficient collection of household plastic waste. *Waste Management* 34(4):719–729.

Daganzo CF (1984) The distance traveled to visit n points with a maximum of c stops per vehicle: An analytic model and an application. *Transportation Science* 18(4):331–350.

Demeester L, Qi M, Van Wassenhove LN (2013) Plant networks for processing recyclable materials. *Manufacturing & Service Operations Management* 15(4):670–688.

Dijkgraaf E, Vollebergh HR (2004) Burn or bury? a social cost comparison of final waste disposal methods. *Ecological Economics* 50(3-4):233–247.

Elmachtoub AN, Grigas P (2022) Smart "predict, then optimize". *Management Science* 68(1):9–26.

Elmachtoub AN, Lam H, Zhang H, Zhao Y (2023) Estimate-then-optimize versus integrated-estimation-optimization versus sample average approximation: A stochastic dominance perspective. *arXiv preprint arXiv:2304.06833* .

EPA (2020) National overview: Facts and figures on materials, waste, and recycling URL `https://www.epa.gov/facts-and-figures-about-materials-waste-and-recycling/national-overview-facts-and-figures-materials#NationalPicture`.

EPA (2021) National recycling strategy: Part one of a series on building a circular economy for all URL `https://www.epa.gov/system/files/documents/2021-11/final-national-recycling-strategy.pdf`.

Fernández-Loría C, Provost F (2022) Causal decision making and causal effect estimation are not the same… and why it matters. *INFORMS Journal on Data Science* 1(1):4–16.

Fleischmann M, Beullens P, Bloemhof-Ruwaard JM, Van Wassenhove LN (2001) The impact of product recovery on logistics network design. *Production and Operations Management* 10(2):156–173.

Folz DH (1995) The economics of municipal recycling: a preliminary analysis. *Public Administration Quarterly* 299–320.

Gambella C, Maggioni F, Vigo D (2019) A stochastic programming model for a tactical solid waste management problem. *European Journal of Operational Research* 273(2):684–694.

Gellynck X, Jacobsen R, Verhelst P (2011) Identifying the key factors in increasing recycling and reducing residual household waste: A case study of the flemish region of belgium. *Journal of Environmental Management* 92(10):2683–2690.

Govindan K, Fattahi M, Keyvanshokooh E (2017) Supply chain network design under uncertainty: A comprehensive review and research directions. *European Journal of Operational Research* 263(1):108–141.

Govindan K, Soleimani H, Kannan D (2015) Reverse logistics and closed-loop supply chain: A comprehensive review to explore the future. *European Journal of Operational Research* 240(3):603–626.

Hemmelmayr VC, Doerner KF, Hartl RF, Vigo D (2014) Models and algorithms for the integrated planning of bin allocation and vehicle routing in solid waste management. *Transportation Science* 48(1):103–120.

Hollmann N, Müller S, Purucker L, Krishnakumar A, Körfer M, Hoo SB, Schirrmeister RT, Hutter F (2025) Accurate predictions on small data with a tabular foundation model. *Nature* 637(8045):319–326.

Janjevic M, Merchán D, Winkenbach M (2021) Designing multi-tier, multi-service-level, and multi-modal last-mile distribution networks for omni-channel operations. *European Journal of Operational Research* 294(3):1059–1077.

Janjevic M, Winkenbach M, Merchán D (2019) Integrating collection-and-delivery points in the strategic design of urban last-mile e-commerce distribution networks. *Transportation Research Part E: Logistics and Transportation Review* 131:37–67.

Jenkins RR, Martinez SA, Palmer K, Podolsky MJ (2003) The determinants of household recycling: a material-specific analysis of recycling program features and unit pricing. *Journal of Environmental Economics and Management* 45(2):294–318.

Kannan R, Bayraksan G, Luedtke JR (2025) Data-driven sample average approximation with covariate information. *Operations Research* .

Kinnaman TC (2006) Policy watch: examining the justification for residential recycling. *Journal of Economic Perspectives* 20(4):219–232.

Kinnaman TC, Fullerton D (2000) Garbage and recycling with endogenous local policy. *Journal of Urban Economics* 48(3):419–442.

Kleywegt AJ, Shapiro A, Homem-de Mello T (2002) The sample average approximation method for stochastic discrete optimization. *SIAM Journal on Optimization* 12(2):479–502.

Laidley TM (2013) The influence of social class and cultural variables on environmental behaviors: Municipal-level evidence from massachusetts. *Environment and Behavior* 45(2):170–197.

MassDEP (2019) Massachusetts materials management capacity study. URL `https://www.mass.gov/files/documents/2019/02/19/capstudy.pdf`.

MassDEP (2021) Massachusetts 2030 solid waste master plan URL `https://www.mass.gov/doc/2030-solid-waste-master-plan-working-together-toward-zero-waste/download`.

MassDEP (2022) Recycling & solid waste data for massachusetts cities & towns. URL `https://www.mass.gov/lists/recycling-solid-waste-data-for-massachusetts-cities-towns`.

MassDEP (2024) 2023 solid waste data update november 2024. URL `https://www.mass.gov/doc/2023-solid-waste-data-update/download`.

McKie EC, Chandrasekaran A, Venkataraman S (2024) How do curbside feedback tactics impact households' recycling performance? evidence from community programs. *Production and Operations Management* 33(5):1064–1082.

Mueller W (2013) The effectiveness of recycling policy options: Waste diversion or just diversions? *Waste Management* 33(3):508–518.

Neal B (2020) *Introduction to Causal Inference*.

NERC (2024) Northeast materials recovery facilities (mrfs) commodity values report. URL `https://www.nerc.org/quarterly-mrf-reports`.

Park S (2018) Factors influencing the recycling rate under the volume-based waste fee system in south korea. *Waste Management* 74:43–51.

Recycle Smart MA (2023) Recycling road map. URL https://recyclesmartma.org/where-does-it-go/.

Reddy KN, Kumar A, Choudhary A, Cheng TE (2022) Multi-period green reverse logistics network design: An improved benders-decomposition-based heuristic approach. *European Journal of Operational Research* 303(2):735–752.

Sadana U, Chenreddy A, Delage E, Forel A, Frejinger E, Vidal T (2024) A survey of contextual optimization methods for decision-making under uncertainty. *European Journal of Operational Research* .

Saragih A, Janjevic M, Winkenbach M (2024a) Revitalizing recycling: Review of the current state, policies, and lessons learned for the future. *MIT Science Policy Review* 5:50–57.

Saragih A, Pontes L, Amin S, Fransoo JC (2024b) Spin the bottle bill: Deposit-refund system policy and reverse supply chain design. *MIT Center for Transportation & Logistics Research Paper Series* .

Sidique SF, Joshi SV, Lupi F (2010) Factors influencing the rate of recycling: An analysis of minnesota counties. *Resources, Conservation and Recycling* 54(4):242–249.

Spinelli A, Maggioni F, Ramos TRP, Barbosa-Póvoa AP, Vigo D (2024) A rolling horizon heuristic approach for a multi-stage stochastic waste collection problem. *European Journal of Operational Research* .

Starr J, Nicolson C (2015) Patterns in trash: Factors driving municipal recycling in massachusetts. *Resources, Conservation and Recycling* 99:7–18.

US Census Bureau (2022a) County population totals and components of change: 2020-2022. URL https://www.census.gov/data/tables/time-series/demo/popest/2020s-counties-total.html.

US Census Bureau (2022b) County subdivision data. URL https://data.census.gov/.

Van Engeland J, Beliën J, De Boeck L, De Jaeger S (2020) Literature review: Strategic network optimization models in waste reverse supply chains. *Omega* 91:102012.

Winkenbach M, Kleindorfer PR, Spinler S (2016) Enabling urban logistics services at la poste through multi-echelon location-routing. *Transportation Science* 50(2):520–540.

Zhu Q, Yu X, Bayraksan G (2024) Residuals-based contextual distributionally robust optimization with decision-dependent uncertainty. *arXiv preprint arXiv:2406.20004* .

E-companion to "Regress, Reverse, Recycle: Contextual Stochastic Optimization in Waste Policy and Logistics Network Design"

EC.1. Proof of Proposition 1

We show the two parts of the proposition separately in the following:

Part 1. By the setup in Section 5.1, the two potential outcomes for municipality i are given by

$$Y_i^0(X_i) = \alpha + X_i^\top \beta + \epsilon_i \quad \text{and} \quad Y_i^1(X_i) = \alpha + X_i^\top \beta + \tau + \epsilon_i.$$

Under the Stable Unit Treatment Value Assumption (SUTVA) and the Modularity Assumption in Section 5.1, the observed outcome is

$$Y_i = (1 - D_i)Y_i^0(X_i) + D_i Y_i^1(X_i) = \alpha + X_i^\top \beta + \tau D_i + \epsilon_i.$$

Taking the conditional expectation given X_i and D_i and invoking the Conditional Independence Assumption in Section 5.1 (which implies $\mathbb{E}[\epsilon_i \mid X_i, D_i] = 0$), we obtain

$$\mathbb{E}[Y_i \mid X_i, D_i] = \alpha + X_i^\top \beta + \tau D_i.$$

Thus, the true conditional mean is identifiable as

$$f^*(X_i, D_i) = \mathbb{E}[Y_i \mid X_i, D_i] = \alpha + X_i^\top \beta + \tau D_i.$$

Part 2. Consider the linear model

$$Y_i = \alpha + X_i^\top \beta + \tau D_i + \epsilon_i,$$

where, by Assumptions in Section 5.1, the error term ϵ_i satisfies $\mathbb{E}[\epsilon_i \mid X_i, D_i] = 0$ and is independent of X_i and D_i.

Define the regressor vector $W_i = \begin{pmatrix} 1 \\ X_i \\ D_i \end{pmatrix}$ and let $\theta = \begin{pmatrix} \alpha \\ \beta \\ \tau \end{pmatrix}$. Then the model can be written compactly as

$$Y_i = W_i^\top \theta + \epsilon_i.$$

The Ordinary Least Squares (OLS) estimator for θ based on a sample $\{(W_i, Y_i)\}_{i=1}^n$ is given by

$$\hat{\theta}_n = \left(\frac{1}{n} \sum_{i=1}^n W_i W_i^\top \right)^{-1} \left(\frac{1}{n} \sum_{i=1}^n W_i Y_i \right).$$

By the Law of Large Numbers, under the assumed regularity conditions, we have:

$$\frac{1}{n}\sum_{i=1}^{n} W_i W_i^\top \xrightarrow{p} \mathbb{E}[W_i W_i^\top] \quad \text{and} \quad \frac{1}{n}\sum_{i=1}^{n} W_i Y_i \xrightarrow{p} \mathbb{E}[W_i Y_i].$$

Since

$$\mathbb{E}[W_i Y_i] = \mathbb{E}[W_i(W_i^\top \theta + \epsilon_i)] = \mathbb{E}[W_i W_i^\top]\theta + \mathbb{E}[W_i \epsilon_i],$$

and $\mathbb{E}[W_i \epsilon_i] = 0$ (by the zero conditional mean of ϵ_i), it follows that

$$\frac{1}{n}\sum_{i=1}^{n} W_i Y_i \xrightarrow{p} \mathbb{E}[W_i W_i^\top]\theta.$$

Thus,

$$\hat{\theta}_n = \left(\frac{1}{n}\sum_{i=1}^{n} W_i W_i^\top\right)^{-1} \left(\frac{1}{n}\sum_{i=1}^{n} W_i Y_i\right) \xrightarrow{p} \left(\mathbb{E}[W_i W_i^\top]\right)^{-1} \mathbb{E}[W_i W_i^\top]\theta = \theta.$$

This implies that

$$\hat{\alpha} \xrightarrow{p} \alpha, \quad \hat{\beta} \xrightarrow{p} \beta, \quad \text{and} \quad \hat{\tau} \xrightarrow{p} \tau,$$

which establishes the consistency of the estimator $\hat{f}_n(x,d) = \hat{\alpha} + x^\top \hat{\beta} + \hat{\tau} d$. $\square$

EC.2. Materials Recycling Facility and Public Authorities Economic Structure

The financial viability of MRF operators depends on balancing revenue with costs:

- Revenue sources include tipping fees paid by municipalities ($90/ton on average; (MassDEP 2022)) and scrap revenue ($124.93/ton; November 2024 data from (NERC 2024)).
- Costs include processing expenses ($87.07/ton; (NERC 2024)), collection operations costs (policy-dependent), fixed facility costs ($385 per ton capacity (Basuhi et al. 2021), 30 year amortization), and disposal costs for residuals ($60.76/ton; approximately 9.1% of incoming recyclables; (MassDEP 2019)).

From the public authorities' perspective:

- Infrastructure costs are $10 per household annually for curbside systems and $4 per household annually for drop-off systems (minimum $20k/year).
- Environmental costs total $170.72/ton due to lost scrap value ($124.93/ton) and social costs from waste disposal ($45.24/ton; parameters from (Dijkgraaf and Vollebergh 2004), given that 86% waste is burned and 14% is landfilled (MassDEP 2024)). MassDEP (2021) targets 30% of MSW to be reduced by 2030.

EC.3. Recycling Network Infrastructure and Operations

The Massachusetts recycling network includes nine operational MRFs whose capacities are shown in Figure EC.1.

Figure EC.1 Capacity of Material Recovery Facilities (MRFs) in MA

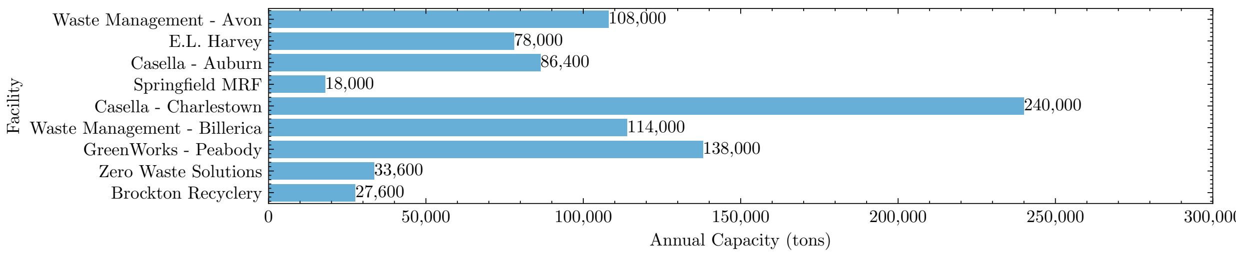

EC.4. Municipal Characteristics Across Recycling Policy

To understand how municipal characteristics differ by optimal policies recommended by CSO solution, we perform the Kruskal-Wallis test with post hoc Dunn analysis and Benjamini-Hochberg correction. This test enables comparison of non-normally distributed data across groups, identifies significant group differences, and controls for multiple comparisons to reduce false positives. The results from the analysis are summarized below and visualized in Figure EC.2, alongside boxplot comparisons. Note that p-value annotations are as follows: ns for $0.05 < p \leqslant 1$, * for $0.001 < p \leqslant 0.05$, ** for $1 \times 10^{-3} < p \leqslant 0.001$, *** for $1 \times 10^{-4} < p \leqslant 1 \times 10^{-3}$, and **** for $p \leqslant 1 \times 10^{-4}$.

Following the framework in Section 4, no significant differences were found between the Curbside and Drop-off policies in terms of income, education, and poverty, but both policies differ significantly from the No Policy, indicating that they tend to be implemented in wealthier, more educated, and less impoverished areas. Socioculturally, Curbside municipalities showed higher foreign-born and non-English speaking populations, and lower white populations compared to Drop-off areas, with both policy types exhibiting these traits compared to municipalities without a policy. Demographically, Curbside areas have younger populations, lower home ownership, and higher density than Drop-off areas, with both showing these characteristics compared to no-policy municipalities. Family size, however, was not significantly different between Curbside and Drop-off, though it differed from no-policy municipalities. This suggests that the optimized solution tends to CSO policies in municipalities with these specific sociocultural and demographic traits, which MA policymakers can take note of in their policymaking.

Comparing with the regression results in Section 7.1, we observe notable characteristics of municipalities with Drop-off and Curbside policies. First, these municipalities tend to have higher education rates, which contribute to higher recycling rates. However, they also typically exhibit higher language diversity, which, interestingly, contributes to lower recycling rates. This highlights that not all characteristics of these municipalities necessarily lead to higher recycling outcomes.

Figure EC.2 **Boxplot Comparison of Municipal Characteristics Across Policies (C: Curbside, D: Drop-off, N: No Policy) via Kruskal-Wallis Test with Post Hoc Dunn Analysis and Benjamini-Hochberg Correction**

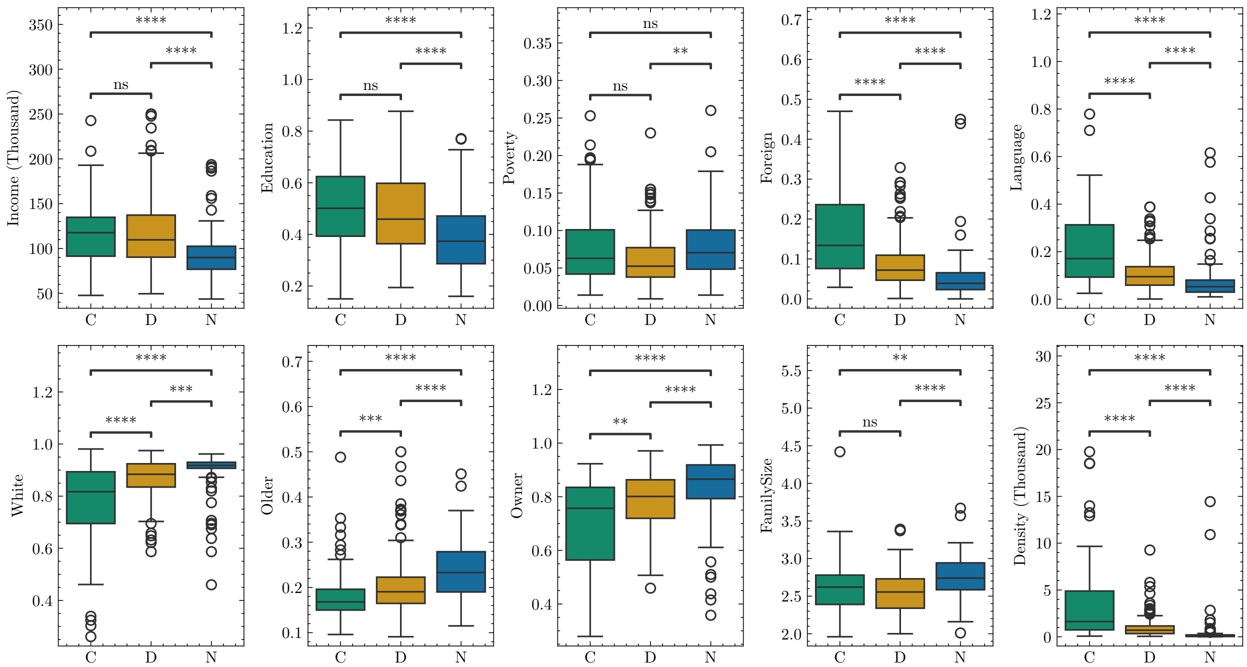

EC.5. Cross-Validated Models Performance

Although our contribution lies in the fact PDS RLasso allows us to manage endogenous uncertainty and policy, we also show that in terms of predictive accuracy, our approach outperforms other prediction models. Figure EC.3 shows the performance comparison of six models using 5-fold and 10-fold cross-validation, measured by R^2 and Mean Absolute Error (MAE). The OLS with Post Double Selection with Rigorous Lasso (PDS RLasso) demonstrates superior performance with the highest R^2 of 0.383 (5-fold) and 0.347 (10-fold), along with consistently low MAE values around 0.053-0.054. Traditional machine learning models like XGBoost ($R^2 = 0.304$ in 5-fold CV) performed notably worse, while LightGBM ($R^2 = 0.377$), TabPFN ($R^2 = 0.381$), and Random Forest ($R^2 = 0.367$) showed intermediate performance in 5-fold CV. The pattern remains consistent in 10-fold CV, though with slightly lower R^2 values across all models. Overall, this suggests that the simpler OLS with PDS RLasso approach is more effective for this prediction task than complex machine learning models and even foundation model by Hollmann et al. (2025).

Figure EC.3 Comparison of Cross-Validated Models' Performance