A Data-Driven County-Level Budget Allocation Model for Opioid Crisis Management: Insights from West Virginia

Mahdi Saadati* Ankit Bansal† Jean-Philippe Richard‡ Samuel Workman§
Gordon Smith¶ Henry Brownstein ‖

Abstract

The opioid crisis has remained a major public health challenge in the United States for many years. This study develops a data-driven decision support framework to guide policymakers in allocating county-level budgets across multiple expenditure categories in order to address the opioid crisis. We compile and curate a detailed dataset on fiscal policy and opioid-related outcomes in West Virginia (WV), the state most severely affected by the epidemic. Drawing on this dataset, we identify causal links between county-level budget allocations and two critical outcomes central to the opioid crisis, number of Opioid-Related Deaths (ORDs) and Crime Incidents (CIs). To capture these relationships, we employ *Tree Ensembles* (TEs) which are trained to predict outcomes as a function of budgetary decisions. We then embed the trained TEs within a Mixed Integer Linear Programming model that produces budget allocation strategies across expenditure categories that maximize the worst-case, risk-averse utility of decision makers across the two outcomes. Our results show that the presented models generate budget allocations that significantly reduce predicted ORDs and CIs across most of the 12 southern WV counties, the area hardest hit by the opioid crisis in the state. For example, in Cabell County, the recommended allocations suggest that predicted levels of both outcomes could have been reduced by approximately 20% in 2023 without increasing the total budget. More broadly, the findings demonstrate that the proposed approach yields tailored, actionable, county-level budget recommendations that can assist policymakers in reprioritizing expenditures to more effectively address the opioid crisis.

Keywords: Opioid crisis; County-level budget allocation; Tree Ensembles; Mixed-integer optimization; West Virginia.

Acknowledgment: This material is based upon work supported by the National Science Foundation under Grant Numbers 2240361, 2240362, 2240363. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

1 Introduction

In recent years, the number of people using, misusing, and abusing heroin and other opioids has reached alarming levels across the United States (U.S.) (Hedegaard et al. 2017, Mattson et al. 2021). This has raised concern among policymakers at both federal and state levels (Office of National Drug Control Policy 2022). Limiting its spread and mitigating its impact are at the forefront of public discourse across

*School of Systems Science and Industrial Engineering, State University of New York-Binghamton, Binghamton, NY, USA
†School of Systems Science and Industrial Engineering, State University of New York-Binghamton, Binghamton, NY, USA, abansal@binghamton.edu (*Corresponding author*)
‡Department of Industrial and Systems Engineering, University of Minnesota, Minneapolis, MN, USA
§Department of Political Science, West Virginia University, Morgantown, WV, USA
¶Department of Epidemiology and Biostatistics, West Virginia University, Morgantown, WV, USA
‖Department of Sociology and Anthropology, West Virginia University, Morgantown, WV, USA

the U.S.. Opioid abuse involves not only heroin, but also modern prescription and synthetic drugs, such as oxycodone and fentanyl (Kolodny et al. 2015), which have legitimate use in surgical procedures and for the treatment of pain (Jones et al. 2018, Marshall et al. 2019). While heroin and today's prescription opioids are pharmacologically similar (Compton et al. 2016), the contemporary opioid epidemic differs from earlier heroin epidemics (Holland et al. 2021) as the people misusing and abusing opioids cross geographical and social boundaries (Mars et al. 2014).

While recent DEA National Threat Assessment reports (Drug Enforcement Administration 2025) provide substantial insight into addressing the opioid epidemic at the national level, there is a critical need for research focused on understanding and evaluating mitigation efforts at the county level. Focusing on counties inside of a state is meaningful since laws and priorities vary significantly across states and resource allocation decisions are easier to coordinate within a state. This paper focuses on the state of West Virginia (`WV`), which has been called the epicenter of the opioid epidemic in the U.S. (Bowden et al. 2018, Merino et al. 2019). `WV` recorded the highest opioid overdose death rate in the U.S. in 2022, with 80.9 deaths per 100,000 population and a total of 1,335 fatalities. Our focus on `WV` allows us to reach the dual goal of studying a state where the need is greatest and whose rural topography allows a more complete view of county-level perspectives often clouded in policy studies examining large metropolitan areas. Further, `WV` has displayed great county-level diversity in tackling the opioid epidemic. Overdoses vary widely across counties as do expenditures on the strategies to mitigate the crisis. Additionally, our focus on counties is justified by the fact that (i) county governments house many governmental functions directly related to supply and demand side crisis interdiction (e.g., emergency, social, and public safety services), and (ii) counties are relatively more important government units in (hardest hit) rural regions as they house and coordinate a disproportionate amount of functions categorized as mitigation strategies for the opioid epidemic.

This paper explores how counties can use the mix of policy tools at their disposal to reduce the impact of the opioid epidemic (Workman and Thomas 2025). Specifically, given evidence of relationship between fiscal decisions and opioid overdose deaths in counties (Lindenfeld et al. 2025), we examine how trade-offs among county-level budget allocations can influence outcomes in the opioid crisis. Our approach differs from most existing research (whether theoretical or applied) that tends to think of public policy as an intervention on the supply-side (e.g., drug seizures) or demand-side (e.g., like harm reduction or prevention). We view the set of policies as a mix of interdependent decisions that might address or mitigate the crisis. Further, these studies rarely focus on rural areas, despite their often disproportionately high overdose rates (Yaemsiri et al. 2019, Gong et al. 2019). We posit that county budget allocations are an important factor to examine because they inherently reflect a county's priorities and policy strategies (Jones and Baumgartner 2005). To test this hypothesis, we use the Arellano-Bond model (Arellano and

Bond 1991, Wooldridge 2010) to estimate the causal impact of expenditure allocations across categories on outcomes associated with the opioid crisis.

While our mathematical framework is applicable to any set of outcomes important to a county, this paper focuses specifically on the reduction of Opioid-Related Deaths (ORDs) and Crime Incidents (CIs). Given the unprecedented severity of the opioid epidemic, these metrics have emerged as primary outcomes of concerns for local governments. They represent the specific areas where decision-makers are most prepared to proactively reallocate and "find room" in their constrained budgets. They capture the duality of the crisis: the direct devastation of individuals and families alongside broader societal consequences like crime. Because these issues fall under core county mandates (including public safety, social services, and emergency response) addressing them is a moral imperative. Furthermore, targeting these outcomes yields secondary benefits, such as reducing high-risk youth behaviors, such as substance abuse and arrests, which remains a top-tier priority for local leadership (Ogunade et al. 2022).

Government budgeting is a complex balancing act, requiring leaders to synthesize constituent demands, peer benchmarks, and a vast array of empirical data. To navigate this complexity, we propose a data-driven optimization model designed to suggest what budget changes (whether through internal reallocations or the strategic use of supplemental state and federal funding) can most effectively mitigate the epidemic. Our model provides a focused analytical tool that is meant to supplement, and not replace, existing decision-making processes by identifying which annual budget expenditure categories may benefit from increased or decreased investment to reduce ORDs and CIs. Beyond its immediate application to the opioid crisis, the underlying framework and its derivations are fully generalizable, offering a robust roadmap for any performance metrics a county may wish to optimize.

Nonetheless, constructing this model posed several challenges because

(a) The necessary data to develop such a model was not readily accessible. For instance, annual aggregated records of budget allocations across different expenditure categories over time for all 55 counties in WV were lacking before this work.

(b) Policymakers care about a range of outcomes in their counties, but their exact priorities for these outcomes are not fully known. Generally, they aim to achieve lower levels of ORDs and CIs. However, the way they assign relative importance to these different outcomes is often unclear.

(c) The outcomes, ORDs and CIs, are complex functions of budget allocations. Such functions are not known explicitly but may be elicited from data. To develop a model that recommends budget allocations to improve outcomes, it is essential to obtain explicit representations of these functions.

To overcome (a), we compile an original data set of local government expenditures for all 55 counties in WV. This data set offers two distinct advantages over the typical Census of Governments data (Lindenfeld

et al. 2025) used in studies of state and local politics. First, it comes directly from mandatory reporting to the WV State Auditor's Office while Census of Governments data is a survey relying on varying response rates, accurate recall and cataloging. Second, our data is annual, reported directly to the state and contains all major expenditure categories. We address challenge (b) by introducing a robust utility optimization framework that aims to maximize the minimum (worst-case) risk-averse utility experienced by decision-makers. In particular, we identify a budget allocation across different expenditure categories that optimizes this worst-case risk-averse utility. To tackle challenge (c), we use machine learning methods trained on historical data to capture the complex relationships between budget allocation decisions and both ORDs and CIs. Specifically, we employ Tree Ensembles (TEs) to predict ORDs and CIs based on budget allocations across different expenditure categories. TEs have proven effective in capturing complex nonlinear relationships between independent variables (in our case budget allocations) and dependent variables (in our case ORDs and CIs). As a result, they have been widely applied in various fields (Ferreira et al. 2016, Deepa et al. 2010, Herrera et al. 2010).We integrate the trained TEs models into a Mixed Integer Linear Programming (MILP) model for optimal budget allocation. Subsequently, we analyze the outputs of the model across the counties in WV that are most severely impacted by the opioid crisis.

Figure 1 presents an overview of the decision support framework we propose. The first step involves obtaining relevant expenditure data from all 55 counties in WV for the years 2012–2023. We then integrate this expenditure data with counts of ORDs and CIs and a range of socio-economic variables. Next, we apply the Arellano-Bond model to identify causal relationships between spending in different expenditure categories, socio-economic variables, and ORDs and CIs. Subsequently, we use the expenditure categories and socio-economic variables for which causal relationships are established to train TEs to predict the number of ORDs and CIs as a function of budget allocations. Lastly, we integrate the trained TEs in an MILP, which is then solved to obtain optimal budget allocations across different expenditure categories. The rest of the paper is structured as follows. Section 2 reviews related literature. Section 3 details our data collection efforts. Section 4 explores the causal relationships between the budget allocations across different expenditure categories and both ORDs and CIs. Section 5 introduces the proposed budget allocation model. Section 6 analyzes the model results for selected counties in WV. Lastly, Section 7 offers concluding remarks. Furthermore, this paper is accompanied by an Online Supplement, where the labels for Sections, Tables, and Figures are prefixed with OS.

2 Related Work

Computational research on the opioid epidemic has focused on modeling progression and dynamics (Luo and Stellato 2024). Many studies adopt the susceptible-infected-recovered (SIR) compartmental model to simulate and forecast the epidemic's trajectory (White and Comiskey 2007, Battista et al. 2019).

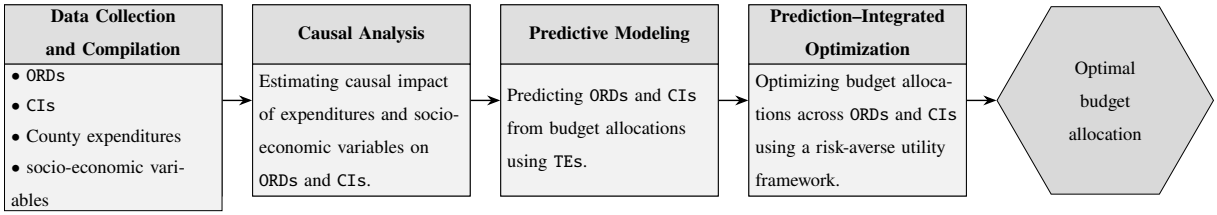

Figure 1: Data-driven decision support framework for county-level budget allocation.

Expanding on this approach, Chen et al. (2019) develop a compartmental model to estimate opioid-related fatalities in the U.S. under both baseline conditions and various intervention scenarios, such as decreasing the misuse of prescription opioids. Pitt et al. (2018) and Rao et al. (2021) use similar models to evaluate outcomes such as life years, quality-adjusted life years, and opioid-related deaths under different policy strategies such as limiting prescription opioid use.

Some studies have employed the SIR compartmental model to optimize resource allocation strategies for addressing the opioid crisis. Ansari et al. (2024) integrate the SIR compartmental model with a Markov Decision Process (MDP) framework to identify optimal budget allocations across preventive and mitigation interventions, with the goal of minimizing the economic cost associated with fatal opioid overdoses. Similarly, Luo and Stellato (2024) combine the SIR model with a Mixed Integer Programming approach to determine both the optimal allocation of opioid treatment facilities and their budgets across each state in the United States. Their objective is to minimize the number of opioid-related overdose fatalities and the population affected by opioid use disorder. Baucum et al. (2025) address the problem of optimally allocating Substance Use Disorder (SUD) treatment centers across counties within a state using a *predict-then-optimize* framework. Rather than relying on the SIR model, they employ a two-stage least squares method to estimate the effects of their interventions. Using the SIR model to assess the impact of policy interventions requires a closed-form expression that links each intervention to the transition rates between compartments. However, for the problem addressed in this paper, establishing such a relationship between county-level budget allocations across different expenditure categories and the transition rates of the SIR model is difficult. Therefore, we propose using TE-based predictive models to estimate ORDs and CIs as functions of county-level budget allocations across various expenditure categories.

Similar to the problem addressed in this paper, there are numerous applications in the literature where the relationship between decisions and outcomes is not explicitly known due to the complexity of the underlying system. However, historical data is available and can be utilized to learn this relationship through a machine learning model, which is then integrated into the decision-making process. Bertsimas et al. (2016) employ ridge regression to predict the efficacy and toxicity outcomes of clinical trials, and incorporate these predictions into a decision-making framework to identify optimal chemotherapy regimens for cancer treatment. Baardman et al. (2019) address a revenue maximization problem where

multiplicative regression is used to estimate demand for promotion vehicles. This work led to efforts to develop mathematical programming formulations of trained machine learning models. Anderson et al. (2020) and Fischetti and Jo (2018) develop MILP formulations that capture the output of a trained neural network for any given input. The resulting formulations can be integrated into optimization models, enabling optimal decision-making where the neural network defines part of the objective or constraints. In a similar vein, Mišić (2020) and Kim et al. (2025) develop MILP formulations to represent the output of trained TEs. In this paper, we use these results to formulate an optimization model that identifies optimal budget allocations across expenditure categories where the outcomes from these allocations are predicted using TEs.

The major contributions of this paper are as follows:

(a) Our effort represents the most comprehensive data collection and curation of local fiscal policy in the state *most* affected by the opioid epidemic. Our data is more detailed and more consistent year-to-year than for studies using the Census of Governments.

(b) We identify causal effects of budget allocations across expenditure categories to both ORDs and CIs using the Arellano–Bond Model (Arellano and Bond 1991, Wooldridge 2010).

(c) We use historical data from Part (a) and results of the Arellano–Bond model from Part (b) to train TEs that capture the relationship between budget allocations and both ORDs and CIs. Leveraging these models, we develop a robust utility optimization framework that identifies budget allocations across expenditure categories to maximize policymakers' worst-case, risk-averse utility in reducing ORDs and CIs.

(d) We conduct an evaluation of the model for 12 counties in southern WV, which are among the hardest hit by the opioid crisis. Our framework provides policy-relevant budget recommendations that are predicted to reduce both ORDs and CIs, and it can be easily adapted to any set of outcomes considered important by the policymakers.

3 Data Collection and Compilation

We focus on two opioid epidemic-related outcomes: ORDs and CIs in WV. For ORDs, we include all drug overdose fatalities associated with opioid use. Detailed toxicological analyses are conducted on all drug-related deaths that occur in the state by the state-wide West Virginia Medical Examiner's Office. All drugs contributing to death are recorded on the death certificate and forwarded to the Bureau for Public Health Office of Vital Statistics who maintain a database of all drug–related overdose deaths for each specific drug detected. Any of the deaths that involved an opioid were included in our analyses as ORDs.

For CIs, following Sim (2023) and Dave et al. (2021), we consider both property and violent crimes. CIs data are collected from the West Virginia Incident-Based Reporting System (WVIBRS), a state-level implementation of the FBI's National Incident-Based Reporting System (NIBRS). We aggregate county-level annual counts of five types of reported incidents: Burglary, Aggravated Assault, Simple Assault, Intimidation Assault, and Robbery.

We also compile detailed data on county-level government spending across expenditure categories. The U.S. Census Bureau's Census of Governments, compiled through survey responses, is a regular source of data on county government revenues and expenditures. Our data set is more comprehensive and detailed as it covers the totality of local government expenditures in the state.

We collect expenditure information from all 55 counties in WV for fiscal years (FY) 2012–2023. County-level expenditure data is available in digital format on the website of the State Auditor's Office's Local Government Division (West Virginia State Auditor 2025). For each county, we collect the PDFs manually and transfer the data to CSV files via an MS Excel macro. Once the data is collected, we cross-walk categories to ensure consistency in expenditure categories across counties and years. We adjust for inflation using the GDP deflator published by the St. Louis Federal Reserve Bank (U.S. Bureau of Economic Analysis 2025) to make expenditure data comparable across time and space. Specifically, we used the deflator for October of each fiscal year, indexed to FY 2023. This enables us to analyze how counties prioritize spending and how local governments shift their budget allocations over time (Workman and Thomas 2025, Workman et al. 2025). The expenditure data include six categories: *Capital Projects (*CP*), Culture and Recreation (*CR*), General Government (*GG*), Health and Sanitation (*HS*), Public Safety (*PS*), and Social Services (*SS*).*

Recognizing and accounting for the potential impact of social conditions on problems related to drugs and crime (Sistani et al. 2023), we also incorporate county-level socio-economic variables into our decision-making framework. All socio-economic data for WV counties from 2012 to 2023 are obtained from the American Community Survey (ACS), available through the U.S. Census Bureau (U.S. Census Bureau 2025). Cesare et al. (2024) identify four key socio-economic categories that may play a significant role in the opioid crisis: Health Care Access, Housing, Employment, and Food Security. Building on Table 1 of their paper, we include seven variables representing these categories, the exact list of which is recorded in Table OS.3 and also appears at the bottom of Table 1. We also incorporate each county's population as an additional input in our model. For ease of exposition, we categorize population alongside the socio-economic variables. The interested reader will find maps illustrating county-level data for budget, ORDs, CIs, and population in Section OS.1A.

4 Arellano–Bond Model for Dynamic Causal Analysis

In this section, we employ the Arellano–Bond model, a Dynamic Panel Data Model (Arellano and Bond 1991, Wooldridge 2010) to estimate the dynamic causal effects of budget allocations and socio-economic variables on ORDs and CIs. This approach is appropriate because our dataset has a panel structure that includes data on budget allocations, ORDs, CIs, and socio-economic variables for all 55 counties in WV over the 12-year span from 2012 to 2023. We specifically adopt the Arellano–Bond model since ORDs and CIs in a given county and year are influenced by their past values. The Arellano–Bond models for ORDs and CIs in county i during year t are given by:

$$\texttt{ORDs}_{it} = \rho^o \texttt{ORDs}_{i,t-1} + \sum_{j \in N} \beta_j^o \tilde{x}_{ij,t-1} + \sum_{j \in S_e} \gamma_j^o h_{ij,t-1} + \alpha_i^o + \varepsilon_{it}^o \tag{1a}$$

$$\texttt{CIs}_{it} = \rho_1^c \texttt{CIs}_{i,t-1} + \rho_2^c \texttt{CIs}_{i,t-2} + \sum_{j \in N} \beta_j^c \tilde{x}_{ij,t-1} + \sum_{j \in S_e} \gamma_j^c h_{ij,t-1} + \alpha_i^c + \varepsilon_{it}^c. \tag{1b}$$

In these models, N denotes the set of expenditure categories, and S_e represents the set of socio-economic variables. The variable $\tilde{x}_{ijt} = log(x_{ijt})$ where x_{ijt} indicates the budget amount allocated by county i to expenditure category j in year t. The variable h_{ijt} represents the value of socio-economic variable j for county i in year t. The terms α_i^o and α_i^c represent the time-invariant county-specific effects in equations (1a) and (1b), respectively, while ε_{it}^o and ε_{it}^c denote the idiosyncratic error terms in the corresponding models. The coefficients $[\rho^o, [\beta_j^o]_{j \in N}, [\gamma_j^o]_{j \in S_e}]$ and $[\rho_1^c, \rho_2^c, [\beta_j^c]_{j \in N}, [\gamma_j^c] j \in S_e]$ capture the dynamic causal effects of the explanatory variables $[\texttt{ORDs}_{i,t-1}, \mathbf{x}_{i,t-1}, \mathbf{h}_{i,t-1}]$ and $[\texttt{CIs}_{i,t-1}, \texttt{CIs}_{i,t-2}, \mathbf{x}_{i,t-1}, \mathbf{h}_{i,t-1}]$ in (1a) and (1b), respectively. The Arellano–Bond model provides consistent estimates of the dependent variables under the *sequential moment restriction* assumption (Wooldridge 2010), which, for (1a) and (1b), is expressed as:

$$\mathbb{E}\left(\texttt{ORDs}_{it} | [\texttt{ORDs}_{i,t_1-1}, \tilde{\boldsymbol{x}}_{i,t_1-1}, \mathbf{h}_{i,t_1-1}]_{t_1=2}^{t}, \alpha_i^o\right) = \mathbb{E}\left(\texttt{ORDs}_{it} | \texttt{ORDs}_{i,t-1}, \tilde{\boldsymbol{x}}_{ij,t-1}, \mathbf{h}_{i,t-1}, \alpha_i^o\right) \tag{2a}$$

$$\mathbb{E}\left(\texttt{CIs}_{it} | [\texttt{CIs}_{i(t_1-1)}, \tilde{\boldsymbol{x}}_{i,t_1-1}, \mathbf{h}_{i,t_1-1}]_{t_1=2}^{t}, \alpha_i^c\right) = \mathbb{E}\left(\texttt{CIs}_{it} | \texttt{CIs}_{i,t-1}, \texttt{CIs}_{i,t-2}, \tilde{\boldsymbol{x}}_{ij,t-1}, \mathbf{h}_{i,t-1}, \alpha_i^c\right). \tag{2b}$$

Restriction (2a) requires that the expected value of $\texttt{ORDs}_{it}$ does not depend on ORDs, budget allocations, or socio-economic variables from years prior to $t-1$, once the corresponding values in year $t-1$ are taken into account. In other words, if we have accounted for ORDs, budget allocations, and socio-economic variables for a county in year $t-1$, then information from earlier years does not provide any additional predictive value for ORDs in year t. Thus, only the most recent year's data (year $t-1$) is relevant for estimating ORDs, and earlier data becomes redundant once year $t-1$ is considered. A similar interpretation applies to condition (2b) for CIs. In the context of this study, it is reasonable to assume that (2a) and (2b) hold, as ORDs and CIs in year t are likely to be strongly influenced by conditions in the years immediately preceding t.

Table 1 displays the output of the Arellano–Bond model obtained using *STATA 19* (StataCorp 2025), when applied to the panel dataset compiled in Section 3. The first column lists the explanatory variables, whereas the second and third columns report the estimated coefficients in (1a) and their corresponding p-values. Similarly, the fourth and fifth columns display the estimated coefficients in (1b) along with their p-values. In (1a) and (1b), we normalize `ORDs`, `CIs` and budget allocations on a per-100,000 resident basis and apply a logarithmic transformation to the budget variables, as their magnitudes are significantly larger than those of `ORDs` and `CIs`. Finally, L1 and L2 denote one-year and two-year lags, respectively. As is typical in econometrics, we log the expenditure data to account for positive skew and long tails in the distribution of expenditures.

Table 1 shows that the coefficients for log(`CR`), log(`CP`), and log(`GG`) in (1a) are non-zero and statistically significant (p-value < 0.1). Similarly, in (1b), the coefficients for log(`PS`), log(`CP`), and log(`GG`) are also non-zero and statistically significant (p-value < 0.1). These results indicate that, there exists a causal relationship between budget allocations to different expenditure categories in year $t-1$ and `ORDs`, `CIs` in year t. Additionally, among the socio-economic variables, the coefficients for *Median Household Income*, *Population*, and *Households receiving public assistance* in (1a) are both non-zero and statistically significant (p-value < 0.1). This indicates a causal relationship between the values of these socio-economic factors in year $t-1$ and `ORDs` in year t. Similarly, in (1b), the coefficients for *Median Household Income*, *Median Gross Rent*, *Uninsured*, and *Population* are likewise non-zero and statistically significant (p-value < 0.1).

Variable	`ORDs`		`CIs`	
	Coeff.	p-value	Coeff.	p-value
L1.`ORDs` per 100,000 people	-0.46^{***}	0.000	–	–
L1.`CIs` per 100,000 people	–	–	-0.073	0.135
L2.`CIs` per 100,000 people	–	–	-0.23^{***}	0.000
L1.`HS` per 100,000 people (log)	-1.00	0.168	3.57	0.334
L1.`PS` per 100,000 people (log)	-0.16	0.633	-15.31^{*}	0.072
L1.`CR` per 100,000 people (log)	0.80^{***}	0.003	5.17	0.457
L1.`SS` per 100,000 people (log)	-0.43	0.254	-2.45	0.450
L1.`CP` per 100,000 people (log)	-0.62^{*}	0.077	-5.23^{**}	0.023
L1.`GG` per 100,000 people (log)	-12.72^{**}	0.027	-82.94	0.206
L1.Renters paying >30% of income	0.17	0.616	2.41	0.591
L1.Homeowners paying >30% of income	0.09	0.788	-7.41	0.520
L1.Median household income	1.76^{***}	0.005	-0.01^{**}	0.015
L1.Median gross rent	-0.001	0.950	-12.84^{*}	0.090
L1.Uninsured	1.25	0.168	40.16^{***}	0.000
L1.Unemployment rate	-0.18	0.942	7.74	0.445
L1.Population	-0.0035^{*}	0.076	0.0578^{**}	0.008
L1.Households receiving public assistance	2.39^{*}	0.090	-3.87	0.750

Significance levels: * p-value < 0.1, ** p-value < 0.05, *** p-value < 0.01.

Table 1: Results of the Arellano–Bond model.

5 Modeling Framework

In this section, we develop a Robust Risk-Averse Utility Optimization Model to determine an optimal allocation of a county's total budget across six categories (HS, PS, CR, SS, CP, and GG) in a given year. The goal is to reduce ORDs and CIs. While our case study and numerical experiments in Section 6 consider specifically these six budget categories and two outcome variables in WV, we next derive a model for the general case involving n budget categories and m outcomes. This general formulation enables broader applicability across states, outcome measures, and budget structures. For simplicity of notation, we omit the county and year indices in the following discussions and use boldface lowercase letters to denote vectors.

We denote the set of expenditure categories by N and the outcomes of interest by M, where $|N| = n$ and $|M| = m$. We let x_i represent the budget allocation to category i. For a given budget allocation $\mathbf{x}$, $\xi_i(\mathbf{x})$ denotes the value of outcome $i \in M$. We denote the vector of all such outcomes (learned from data) by $\bar{\xi}(\mathbf{x}) = \{\xi_1(\mathbf{x}), \xi_2(\mathbf{x}), \ldots, \xi_m(\mathbf{x})\}$. Using this notation, the model we solve is

$$\max_{\mathbf{x}\in\mathbf{X}} \min_{u\in\mathbf{U}} u(\bar{\xi}(\mathbf{x})), \tag{3}$$

where $u(\cdot) : \mathbb{R}^m \to \mathbb{R}$ is the policymakers' overall utility from a given outcome vector $\bar{\xi}(\mathbf{x})$, where $\mathbf{U}$ is a set of utility functions over the outcomes, and where $\mathbf{X}$ is the set of feasible budget allocations

$$\mathbf{X} = \left\{ \mathbf{x} \in \mathbb{R}^n_+ \;\middle|\; \textstyle\sum_{k\in N} x_k \le b, \quad \hat{x}_k - \kappa^-_k \le x_k \le \hat{x}_k + \kappa^+_k, \quad \forall k \in N \right\}. \tag{4}$$

Model (3) seeks a budget allocation that maximizes the worst-case utility over all utility functions in the set $\mathbf{U}$, which represent plausible preferences policymakers may have regarding reductions in undesirable outcomes. This formulation offers the advantage of a robust approach: it identifies an allocation that performs well even under the least favorable utility function within $\mathbf{U}$. The set $\mathbf{X}$ of feasible budget allocations defined in (4) ensures that: (i) the total available budget b is not over-allocated across n categories, and (ii) the amount allocated to each category k can decrease from the *reference budget allocation* $\hat{x}_k$ by at most κ^-_k or increase by at most κ^+_k. These latter constraints, which we refer to as *reference budget allocation* (RBA) constraints and often write as $\mathbf{x} \in [\hat{\mathbf{x}} - \boldsymbol{\kappa}^-, \hat{\mathbf{x}} + \boldsymbol{\kappa}^+]$, restrict annual changes by keeping each expenditure category's allocation close to the previous year's value, reflecting counties' limited flexibility in adjusting budgets.

Although model (3) is notationally simple, solving it requires an explicit reformulation suitable for implementation using optimization solvers. To obtain this reformulation, we proceed in two steps. First, in Section 5.1, we focus on modeling how budget allocations affect different outcomes. For this task, we make use of TEs, a class of machine learning models known for their strong predictive accuracy and ability to capture complex relationships in diverse datasets. We train these models on the historical

data and socio-economic variables compiled in Section 3, taking into account the causal insights from Section 4. In addition to their predictive accuracy, a key advantage of TEs is that their input-to-output relationships can be represented through mixed-integer linear constraints. This allows us to incorporate the learned vector-function $\bar{\xi}(\mathbf{x})$ directly into (3). Second, in Section 5.2, we describe a set of reasonable assumptions on the set of utility functions $\mathbf{U}$ under which we can reformulate the inner minimization problem over utility functions $u \in \mathbf{U}$ using linear programming duality. By introducing suitable dual variables and constraints, we show that the robust objective can also be captured through linear constraints. Together, these two steps establish that (3) can be formulated and solved as an MILP, which we describe in Section 5.3.

5.1 Modeling Outcome Functions $\xi_i(\mathbf{x})$ using TEs

Tree ensembles are powerful machine learning models that combine the predictions of multiple decision trees to produce more accurate and robust results for both regression and classification. Each decision tree is a piecewise constant function whose structure can be visually represented as a tree. We refer the unfamiliar reader to Hastie et al. (2009) for a textbook discussion or Section OS.1B for a brief introduction.

5.1.1 Pruning Decision Trees

With a slight abuse of notation, we use t to denote both the graphical structure of a decision tree and the function it defines. Specifically, we write $t(\mathbf{x})$ to indicate the output of the tree on input vector $\mathbf{x}$. Decision trees are naturally constructed over hyper-rectangular domains $\mathcal{X} = [\boldsymbol{l}, \boldsymbol{u}]$. In the ensuing sections, we will often consider a decision tree's outputs over a smaller hyper-rectangle $\mathcal{X}'$. Next, we show that such a restriction results in a function that can still be represented by a decision tree. The idea behind this simplification is formalized in Observation 2, which builds on the following observation that any branch of a decision tree that is never reached by inputs from the subdomain $\mathcal{X}'$ can be safely *pruned*, whereas any branch that is always taken by such inputs when the parent node is reached can be *contracted*.

Observation 1. *Let t be a decision tree defined over $\mathcal{X} \subseteq \mathbb{R}^p$, and suppose that the query at node v of t is of the form $x_k \leq d$ versus $x_k > d$, for some $k \in \{1, \ldots, p\}$ and threshold $d \in \mathbb{R}$. Let t_0 denote the subtree of t rooted at v, and let t_1 and t_2 be the left and right subtrees of v, respectively. For bounds $r_k \leq s_k \in \mathbb{R}$, define the restricted domain $\mathcal{X}' = \{x \in \mathcal{X} \mid r_k \leq x_k \leq s_k\}$. Construct a modified tree t' as follows: (i) if $d < r_k$, replace t_0 in t with t_2, (ii) if $d \geq s_k$, replace t_0 in t with t_1, (iii) otherwise, let $t' = t$. Then, for all $x \in \mathcal{X}'$, we have $t(x) = t'(x)$.*

In the special case where $r_k = s_k = \hat{x}_k$, *i.e.*, the value of x_k is known, then one of the condition (i) or (ii) in Observation 1 will be always satisfied. Consequently, the above construction yields a simplified tree t', equivalent to t over the restricted domain, that bypasses the use of node v and eliminates the associated

query on x_k. The result of Observation 1 can be applied recursively to prune all splits involving a variable x_k whose thresholds fall outside the interval $[r_k, s_k)$. In particular, any split of the form $x_k \leq d$ with $d < r_k$, or $x_k > d$ with $d > s_k$, can be removed by replacing the split node with the child corresponding to the consistent branch. This recursive procedure provides an efficient method to build the pruned tree t', and leads to the following observation.

Observation 2. *Let t be a decision tree defined over $\mathcal{X} \subseteq \mathbb{R}^p$. For bounds $r_k \leq s_k \in \mathbb{R}$ for $k \in \{1, \ldots, p\}$, define the restricted domain $\mathcal{X}' = \{x \in \mathcal{X} \mid r_k \leq x_k \leq s_k, \forall k = 1, \ldots, p\}$. There exists a decision tree t' that only uses queries of the form $x_k \leq d$ or $x_k > d$ where $d \in [r_k, s_k)$ for each k such that $t'(x) = t(x)$, $\forall x \in \mathcal{X}'$. Further, if $r_k = s_k$ for some k, t' does not contain any queries on x_k.*

Given a restricted domain $\mathcal{X}'$ and an original decision tree t, we refer to the tree t' satisfying the conditions of Observation 2 as the *restriction* of t over $\mathcal{X}'$.

5.1.2 Formulating Input-to-Ouput Relationships of `TEs` as MIPs

Given their strong predictive performance, we use `TEs` to model outcomes $\chi_i(\mathbf{x}, \mathbf{h})$ for each $i \in M$ as functions of both budget allocations $\mathbf{x}$ and socio-economic variables $\mathbf{h}$. Although, unlike budget allocations, socio-economic variables cannot be altered, they are included during training to improve predictive accuracy. Because (3) is solved for each county individually, the socio-economic variables $\bar{\boldsymbol{h}}$ of a county will be known at the time of solving. This allows us to define the outcome functions solely in terms of $\mathbf{x}$ by setting $\xi_i(\mathbf{x}) = \chi_i(\mathbf{x}, \bar{\boldsymbol{h}})$ for $i \in M$. Observation 2 establishes that $\xi_i(\mathbf{x})$ is described as a `TE` in which all decision nodes query $\mathbf{x}$ variables exclusively. For each outcome $i \in M$, we denote the set of all trees t in the trained `TE` that have been pruned so as to predict $\xi_i(\mathbf{x})$ solely from $\mathbf{x}$ by $\mathcal{T}_i$. For short, we refer to the collection of the set of trees $\mathcal{T}_i$, for all $i \in M$, as $\mathcal{T}$. Given the training data, it is natural to assume that these trees are defined over the non-negative orthant.

We are interested in modeling the relationship between $\mathbf{x}$ and ξ_is over the set of possible budget allocations, and in particular, are interested in deriving an explicit formulation for the set

$$O(\mathcal{T}, \mathbf{X}) = \{(\mathbf{x}, \mathbf{z}) \mid z_i = \xi_i(\mathbf{x}), \forall i \in M, \mathbf{x} \in \mathbf{X}\}.$$

Next, we show that $O(\mathcal{T}, X)$ can be described as a mixed-integer programming set. This description extends the formulation in Kim et al. (2025), which itself built on the earlier work of Mišić (2020).

Let `f`(t) denote the set of leaves of tree $t \in \mathcal{T}$ and L_k denote the collection of all split values involving expenditure category k that arise in any of the trees in $\mathcal{T}$. Specifically, for each $k \in N$, we assume wlog that $L_k = \{a_{k0}, a_{k1}, a_{k2}, a_{k3}, \ldots, a_{k|L_k|}\}$ where $0 = a_{k0} < a_{k1} < \cdots < a_{k|L_k|} = \infty$ and for each $v \in \{1, 2, 3 \ldots, |L_k| - 1\}$ there exists a tree in $t \in \mathcal{T}$ that includes the split condition $x_k \leq a_{kv}$.

The elements of L_k divide the domain of the expenditure category k into $|L_k|$ intervals of the form $(a_{k(v-1)}, a_{kv}]$ for $v \in \bar{L}_k := \{1, 2, \ldots, |L_k|\}$.

Consider a tree $t \in \mathcal{T}_i$. Each leaf of t corresponds to a region in the input space $\mathbf{x}$ that is characterized by a collection of intervals on each of the input variables. For each variable x_k and each leaf l of tree t, we denote the set of these intervals by $J_{lt}^k \subseteq \bar{L}_k$. Finally, we let c_{lt} be the output value of leaf l of tree t and v_t be the weight assigned to tree t. The combination of sets $\mathcal{T}_i$, J_{lt}^k, L_k, and parameters c_{lt} and v_t, is sufficient to completely define and compute all values of all outcomes $\xi_i(\mathbf{x})$ for any combination of input values $\mathbf{x}$.

In addition to the natural variables $\mathbf{z}$ and $\mathbf{x}$, the formulation we use requires the introduction of two sets of binary variables. Let q_{lt} be a binary variable that takes the value 1 if leaf l of tree t is selected and 0, otherwise. Additionally, let $\bar{x}_{kv}$ be a binary variable that takes the value 1 if x_k belongs to the interval $(a_{k(v-1)}, a_{kv}]$ and 0, otherwise. Building on Kim et al. (2025) and Mistry et al. (2021) and using the fact that expenditures are numerical, we write that $O(\mathcal{T}, \mathbf{X}) = \text{proj}_{(\mathbf{x},\mathbf{z})}\, \mathcal{P}(\mathcal{T}, \mathbf{X}) := \{(\mathbf{x}, \bar{\mathbf{x}}, \mathbf{z}, \mathbf{q}) \mid (5a) - (5h), \mathbf{x} \in \mathbf{X}\}$ whose feasible region contains the constraints

$$z_i = \sum_{t \in \mathcal{T}_i} \sum_{l \in \mathtt{f}(t)} c_{lt} v_t q_{lt} \qquad \forall i \in M \tag{5a}$$

$$x_k \geq \sum_{v \in \bar{L}_k} a_{k(v-1)} \bar{x}_{kv} \qquad \forall k \in N \tag{5b}$$

$$x_k \leq \sum_{v \in \bar{L}_k} a_{kv} \bar{x}_{kv} \qquad \forall k \in N \tag{5c}$$

$$\sum_{v \in \bar{L}_k} \bar{x}_{kv} = 1 \qquad \forall k \in N \tag{5d}$$

$$\sum_{l \in \mathtt{f}(t): J_{lt}^k \subseteq \{j_1, \ldots, j_2\}} q_{lt} \leq \sum_{v=j_1}^{j_2} \bar{x}_{kv} \qquad \forall t \in \mathcal{T}_i, \forall i \in M, \forall k \in N, \forall j_1 \leq j_2 \in \bar{L}_k \tag{5e}$$

$$\sum_{l \in \mathtt{f}(t)} q_{lt} = 1 \qquad \forall t \in \mathcal{T}_i, \forall i \in M \tag{5f}$$

$$q_{\cdot t} \in \{0, 1\}^{|\mathtt{f}(t)|} \qquad \forall t \in \mathcal{T}_i, \forall i \in M \tag{5g}$$

$$\bar{\mathbf{x}}_k \in \{0, 1\}^{|\bar{L}_k|} \qquad \forall k \in N. \tag{5h}$$

Constraints (5b)-(5d) ensure that, for each $k \in N$, an interval in $\bar{L}_k$ containing x_k is selected, whereas constraints (5e) and (5f) determine the leaf of each tree that is reached for a given $\bar{\mathbf{x}}_k$. Constraints (5a) aggregate the values of the selected leaves to compute $z_i = \xi_i(\mathbf{x})$. In the literature, the simplex constraint (5d) is often reformulated in incremental form to better leverage the branching capabilities of commercial solvers. Our models solve efficiently without this transformation.

The definition of $\mathbf{X}$, and therefore of $O(\mathcal{T}, \mathbf{X})$, incorporates `RBA` constraints (4). These constraints require that the function described by each `TE` be optimized over a subregion of its domain. Observation 2

shows that these constraints can be enforced implicitly by restricting each tree to the subdomain defined by $[\hat{\mathbf{x}} - \boldsymbol{\kappa}^-, \hat{\mathbf{x}} + \boldsymbol{\kappa}^+]$. As a result, we obtain a simplified representation of the input-to-output relationship between budget allocations and outcomes that does not require explicitly formulating or enforcing the RBA constraints. Proposition 1 follows directly.

Proposition 1. *Define* $\mathcal{X} = [\hat{\mathbf{x}} - \boldsymbol{\kappa}^-, \hat{\mathbf{x}} + \boldsymbol{\kappa}^+]$ *and* $\mathbf{X}' = \{\mathbf{x} \in \mathbb{R}^n_+ \mid \sum_{k \in N} x_k \leq b\}$. *It holds that*

$$\mathcal{P}(\mathcal{T}, \mathbf{X}) = \mathcal{P}(\mathcal{T}', \mathbf{X}'),$$

where $\mathcal{T}'$ *is the collection of trees obtained by restricting the domains of trees of* $\mathcal{T}$ *to* $\mathcal{X}$.

Restricting the trees to the domain $\mathcal{X}$ alters the sets J^k_{lt}, $\mathtt{f}(t)$, and L_k that define their structure. Specifically, after applying this restriction, we have $a_{k0} = \hat{x}_k - \kappa^-_k$ and $a_{k|L_k|} = \hat{x}_k + \kappa^+_k$ for each $k \in N$. Nevertheless, the structural form of the formulation in (5) is preserved following this restriction.

Next, we observe that, although the variables $\mathbf{x}$ were introduced as the natural inputs to the model, the predicted values $\mathbf{z}$ depend only on variables $\bar{\mathbf{x}}$, as evidenced by constraint (5a). We show next that these variables are not needed and can be removed from the formulation. Beyond simplifying the formulation, working with the reduced variable set $\bar{\mathbf{x}}$ highlights a structural property of TE functions: since these functions are piecewise constant, their optimal solutions are not isolated points but entire hyperrectangles of input values yielding the same objective. Thus, instead of identifying a single optimal solution, the reduced model focuses on producing a hyperrectangle of optimal solutions. This has practical advantages, as it provides decision-makers with a range of flexible alternatives rather than a single prescribed choice. We revisit this point during our discussion of the computational results in Section 6. Eliminating variables $\mathbf{x}$ also has the advantage of reducing the number of variables and constraints in the model, without compromising the quality of the LP relaxation of the model. More precisely, we next show that the $\mathbf{x}$ variables used in the definition of $\mathcal{P}(\mathcal{T}', \mathbf{X}')$ can be eliminated from this formulation without weakening its relaxation. The set of constraints from $\mathcal{P}(\mathcal{T}', \mathbf{X}')$ containing only $\mathbf{x}$ define the set:

$$\mathcal{P}_1 = \left\{ (\mathbf{x}, \bar{\mathbf{x}}) \in \mathbb{R}^{|N|} \times [0,1]^{\sum_{k \in N} |\bar{L}_k|} \;\middle|\; \begin{array}{ll} \sum_{k \in N} x_k \leq b, & \\ x_k \geq \sum_{v \in \bar{L}_k} a_{k(v-1)} \bar{x}_{kv}, & \forall k \in N \\ x_k \leq \sum_{v \in \bar{L}_k} a_{kv} \bar{x}_{kv}, & \forall k \in N \end{array} \right\}. \tag{6}$$

The following proposition characterizes the projection of $\mathcal{P}_1$ onto the space of variables $\bar{\mathbf{x}}$.

Proposition 2. *Consider the inequality*

$$\sum_{k \in N} \sum_{v \in \bar{L}_k} a_{k(v-1)} \bar{x}_{kv} \leq b. \tag{7}$$

It holds that $\mathrm{proj}_{\bar{\mathbf{x}}} \mathcal{P}_1 = \mathcal{P}_2$ *where* $\mathcal{P}_2 = \left\{ \bar{\mathbf{x}} \in [0,1]^{\sum_{k \in N} |\bar{L}_k|} \mid (7) \right\}$.

Proof. First, we prove that $\text{proj}_{\bar{\mathbf{x}}} \mathcal{P}_1 \subseteq \mathcal{P}_2$. Let $\bar{\mathbf{x}}^1 \in \text{proj}_{\bar{\mathbf{x}}} \mathcal{P}_1$. This implies that there exists $\mathbf{x}^1$ such that $(\mathbf{x}^1, \bar{\mathbf{x}}^1) \in \mathcal{P}_2$. Using the first two families of constraints of $\mathcal{P}_1$, we write that $b \geq \sum_{k \in N} x^1_k \geq \sum_{k \in N} \sum_{v \in \bar{L}^p_k} a^p_{k(v-1)} \bar{x}^1_{kv}$, which along with the bounds on variables $\bar{\mathbf{x}}$, implies that $\text{proj}_{\bar{\mathbf{x}}} \mathcal{P}_1 \subseteq \mathcal{P}_2$. Second, we prove that $\mathcal{P}_2 \subseteq \text{proj}_{\bar{\mathbf{x}}} \mathcal{P}_1$. Let $\bar{\mathbf{x}}^1 \in \mathcal{P}_2$. Hence, $\bar{\mathbf{x}}^1$ satisfies the bound constraints. Further, define $x^1_k = \sum_{v \in \bar{L}_k} a_{k(v-1)} \bar{x}^1_{kv}$ for $k \in N$. Since $\sum_{k \in N} \sum_{v \in \bar{L}_k} a_{k(v-1)} \bar{x}^1_{kv} \leq b$, it follows that $\sum_{k \in N} x^1_k \leq b$. Furthermore, we have that $x^1_k = \sum_{v \in \bar{L}_k} a_{k(v-1)} \bar{x}^1_{kv} < \sum_{v \in \bar{L}_k} a_{kv} \bar{x}^1_{kv}$, where the last inequality is due to the facts that $a_{k(v-1)} < a_{kv}$ and $\bar{x}^1_{kv} \geq 0$ for all k and v. This implies that $(\mathbf{x}^1, \bar{\mathbf{x}}^1) \in \mathcal{P}_1$, showing that $\mathcal{P}_2 \subseteq \text{proj}_{\bar{\mathbf{x}}} \mathcal{P}_1$. □

Proposition 2 yields an alternative, more compact formulation for representing the outcomes ξ_i as functions of the budget allocation decisions, without loss of tightness. This reformulation can be expressed as:

$$Q(\mathcal{T}') := \{(\bar{\mathbf{x}}, \mathbf{z}, \mathbf{q}) \mid (5a), (5d) - (5h), (7)\}.$$

Hence, the derivations in this section provide the following alternate formulations of (3):

$$\max_{\mathbf{x} \in \mathbf{X}} \min_{u \in \mathbf{U}} u(\bar{\xi}(\mathbf{x})) = \max_{(\mathbf{x}, \bar{\mathbf{x}}, \mathbf{z}, \mathbf{q}) \in \mathcal{P}(\mathcal{T}, X)} \min_{u \in \mathbf{U}} u(\mathbf{z}) = \max_{(\bar{\mathbf{x}}, \mathbf{z}, \mathbf{q}) \in Q(\mathcal{T}')} \min_{u \in \mathbf{U}} u(\mathbf{z}). \tag{8}$$

In the rest of this paper, we will use the latter one.

5.2 Robust Risk-Averse Utility Optimization Model

We now turn our attention to the inner problem of (3), $\min_{u \in \mathbf{U}} u(\mathbf{z})$, which seeks to identify the worst-case utility function within the admissible set $\mathbf{U}$ for a given vector of outcomes $\mathbf{z}$. We begin by motivating this formulation and outlining the assumptions we place on the utility functions in $\mathbf{U}$. We then derive an equivalent reformulation of the problem that can be integrated into (8).

Policymakers often have limited attention, evaluate allocations on multiple outcomes and differ in the weights they attach to the relative importance of outcomes (Jones 2003, Workman 2015). Thus, comparing two vectors of outcomes is difficult. Therefore, we assume policymakers face uncertainty about how their strategies map onto outcomes and their own utility over the outcomes of their investments. Policymakers aim to act rationally but are constrained by the context in which decisions are made. As it is rational to allocate budgets so that utility is maximized (Mas-Colell et al. 1995), we adopt a utility optimization model, maximizing the worst-case utility of the policymakers, for allocating budgets across the expenditure categories with a focus on managing the opioid crisis. Our modeling framework is based on the following assumptions.

Assumption 1.

a. *For a vector* $\mathbf{z}$ *of outcomes* ξ*, the utility function of the policymaker satisfies* $u(\mathbf{z}) = \sum_{i \in M} \sigma_i \bar{u}_i(z_i)$ *where* $\bar{u}_i(z_i) : \mathbb{R} \to [0, 1]$ *is the utility of the policymaker when output* ξ_i *takes value* z_i *and* σ_i *is the priority coefficient for outcome* i *where a higher value of* σ_i *indicates greater importance on improving outcome* i*.*

b. *Let* $\bar{\mathbf{U}}_i$ *be the set of non-decreasing concave piecewise linear functions* $\hat{u}_i(\cdot)$ *that are such that* $\bar{u}_i^a(\cdot) \leq \hat{u}_i(\cdot) \leq \bar{u}_i^b(\cdot)$ *and right-derivatives satisfy* $\bar{u}_i'^a(\cdot) \leq \hat{u}_i'(\cdot) \leq \bar{u}_i'^b(\cdot)$ *for two given reference functions* $\bar{u}_i^a$ *and* $\bar{u}_i^b$*. We assume that* $\bar{u}_i(\cdot) \in \bar{\mathbf{U}}_i \ \ \forall i \in M$*.*

Assumptions 1a and 1b ensure that policymakers are risk-averse (Mas-Colell et al. 1995). This is reasonable considering the state of the opioid crisis in `WV` where policymakers are likely to favor budget allocations that offer modest but reliable improvements over those that promise greater gains with high uncertainty (Kahneman and Tversky 1979). These assumptions also ensure that policymakers' utility is non-decreasing in the outcomes, preferring higher values for each outcome. As lower `ORDs` and `CIs` are preferred, we will transform them to conform with these assumptions.

Assumption 1a allows us to simplify both the modeling and computation by ensuring that the utility from a given outcome $i \in M$ is independent of the values of other outcomes $j \in M$. The piecewise linear assumption in Assumption 1b is justified because piecewise linear functions can closely approximate a wide range of functions. Furthermore, the *reference functions* specified in the same assumption limit the set of functions included in $\bar{\mathbf{U}}_i$, which restricts our framework from producing overly conservative solutions. Let $\mathbf{U}$ denote the set of all utility functions over $\mathbb{R}^m$ that satisfy Assumption 1. Under this definition, the inner minimization problem in (3) reduces to

$$\min_{u \in \mathbf{U}} u(\mathbf{z}) = \sum_{i \in M} \sigma_i \bar{\pi}_i(\mathbf{z}) \quad \text{where} \quad \bar{\pi}_i(\mathbf{z}) = \min_{\bar{u}_i \in \bar{\mathbf{U}}_i} \bar{u}_i(z_i). \tag{9}$$

For an outcome $z_i \in [\theta_i^1, \theta_i^2]$, let $\theta_i^1 = p_{i0} < p_{i1} < p_{i2} < \cdots < p_{iw_i} = \theta_i^2$ represent a partition of $[\theta_i^1, \theta_i^2]$ into w_i intervals. While multiple approaches could be used to define these intervals, we extract them from the trained `TE` as described in Section 6. Given these partitions and Assumption 1, we approximate $\bar{\pi}_i(\mathbf{x})$ by solving the following linear program (Hu and Mehrotra 2015):

$$\bar{\pi}_i(\mathbf{z}) = \min \quad \tilde{b}_i z_i + \tilde{a}_i \tag{10a}$$

s.t.

$$[\hat{\boldsymbol{y}}^1] \qquad \hat{b}_{ij+1} - \hat{b}_{ij} - \dot{b}_{ij}\left(p_{ij+1} - p_{ij}\right) \leq 0, \quad \forall j = 0, \ldots, w_i - 1 \tag{10b}$$

$$[\hat{\boldsymbol{y}}^2] \qquad \hat{b}_{ij+1} - \hat{b}_{ij} - \dot{b}_{ij+1}\left(p_{ij+1} - p_{ij}\right) \geq 0, \quad \forall j = 0, \ldots, w_i - 2 \tag{10c}$$

$$[\hat{\boldsymbol{y}}^3] \qquad l_{ij}\tilde{b}_i + \tilde{a}_i \geq \hat{b}_{ij}, \quad \forall j = 0, \ldots, w_i \tag{10d}$$

$$[\hat{\boldsymbol{y}}^4], [\hat{\boldsymbol{y}}^5] \qquad \tilde{u}_i(p_{ij}, 0.41) \leq \hat{b}_{ij} \leq \tilde{u}_i(p_{ij}, 0.68), \quad \forall j = 0, \ldots, w_i \tag{10e}$$

$$[\hat{\mathbf{y}}^6], [\hat{\mathbf{y}}^7] \qquad \tilde{u}_i'(p_{ij}, 0.41) \leq \dot{b}_{ij} \leq \tilde{u}_i'(p_{ij}, 0.68), \quad \forall j = 0, \ldots, w_i - 1 \tag{10f}$$

$$\hat{b}_{ij}, \tilde{a}_i, \tilde{b}_i, \dot{b}_{ij} \geq 0, \qquad \forall j = 0, \ldots, w_i \tag{10g}$$

where the notation for dual variables is stated before each constraint, and

$$\tilde{u}_i(\hat{t}, \lambda) = \frac{1}{1-\lambda}\hat{t}^{1-\lambda}, \quad \lambda \in [0.41, 0.68] \tag{11}$$

is the *Constant Relative Risk Aversion (CRRA)* utility function (Holt and Laury 2002). We use *CRRA* in our analysis as this utility function has been empirically shown to effectively capture the risk aversion behavior of decision-makers (Holt and Laury 2002, Hu and Mehrotra 2015). Variables $\tilde{b}_i$ and $\tilde{a}_i$ represent the slope and intercept of the linear segment of the piecewise-linear function at z_i. Variable $\hat{b}_{ij}$ captures the value of the function $\bar{\pi}_i$ at the point p_{ij}, whereas variable $\dot{b}_{ij}$ represents the slope of the line segment of the piecewise-linear function $\bar{\pi}_i$ over the interval $[p_{ij}, p_{i,j+1}]$. Constraints (10b) and (10c) ensure that $\bar{\pi}_i$ is piecewise-linear, concave, and non-decreasing. Constraints (10d) along with the objective function (10a) give the value of $\bar{\pi}_i(\mathbf{z})$. Constraints (10e) and (10f) impose bounds on the value of $\bar{\pi}_i$ at the $w_i + 1$ breakpoints and on the slope of each of the w_i line segments, respectively. These constraints approximate the bounds defined by the *reference functions* in Assumption 1b by specifying $\bar{u}^a(\cdot) = \tilde{u}(\cdot, 0.41)$ and $\bar{u}^b(\cdot) = \tilde{u}(\cdot, 0.68)$.

Next, we dualize (10) for each $i \in M$ to obtain an expression for $\bar{\pi}_i(\mathbf{z})$ as a maximization problem. To simplify the expression of this dual problem, we introduce the notation $\hat{\mathbf{y}}_i^{1:7}$ for the vector $(\hat{\mathbf{y}}_i^1, \hat{\mathbf{y}}_i^2, \hat{\mathbf{y}}_i^3, \hat{\mathbf{y}}_i^4, \hat{\mathbf{y}}_i^5, \hat{\mathbf{y}}_i^6, \hat{\mathbf{y}}_i^7)$ of all dual variables of (10), where $\hat{\mathbf{y}}_i^1, \hat{\mathbf{y}}_i^6, \hat{\mathbf{y}}_i^7 \in \mathbb{R}^{w_i}$, $\hat{\mathbf{y}}_i^2 \in \mathbb{R}^{w_i - 1}$, and $\hat{\mathbf{y}}_i^3, \hat{\mathbf{y}}_i^4, \hat{\mathbf{y}}_i^5 \in \mathbb{R}^{w_i + 1}$. For the same reason, we define the function

$$\Psi_i(\hat{\mathbf{y}}_i^{1:7}) = \sum_{j=0}^{w_i} \Big(\tilde{u}_i(p_{ij}, 0.41)\hat{y}_{ij}^4 + \tilde{u}_i(p_{ij}, 0.68)\hat{y}_{ij}^5 + \tilde{u}_i'(p_{ij}, 0.41)\hat{y}_{ij}^6 + \tilde{u}_i'(p_{ij}, 0.68)\hat{y}_{ij}^7 \Big).$$

Using this notation, it is easily established that the dual of (10) is the optimization problem

$$\bar{\pi}_i(\mathbf{z}) = \max \left\{ \Psi_i(\hat{\mathbf{y}}_i^{1:7}) \;\middle|\; (\hat{\mathbf{y}}_i^1, \hat{\mathbf{y}}_i^2, \hat{\mathbf{y}}_i^3, \hat{\mathbf{y}}_i^4, \hat{\mathbf{y}}_i^5, \hat{\mathbf{y}}_i^6, \hat{\mathbf{y}}_i^7, z_i) \in \mathcal{R}_i \right\} \tag{12}$$

whose feasible region

$$\mathcal{R}_i = \left\{ (\hat{\mathbf{y}}_i^1, \hat{\mathbf{y}}_i^2, \hat{\mathbf{y}}_i^3, \hat{\mathbf{y}}_i^4, \hat{\mathbf{y}}_i^5, \hat{\mathbf{y}}_i^6, \hat{\mathbf{y}}_i^7, z_i) \;\middle|\; \text{(14a)} - \text{(14j)} \right\} \tag{13}$$

is defined through the constraints

$$-\hat{y}_{i0}^1 - \hat{y}_{i0}^2 - \hat{y}_{i0}^3 + \hat{y}_{i0}^4 + \hat{y}_{i0}^5 \leq 0 \tag{14a}$$

$$\hat{y}_{i(j-1)}^1 - \hat{y}_{ij}^1 + \hat{y}_{i(j-1)}^2 - \hat{y}_{ij}^2 - \hat{y}_{ij}^3 + \hat{y}_{ij}^4 + \hat{y}_{ij}^5 \leq 0 \qquad \forall j = 1, \ldots, w_i - 2 \tag{14b}$$

$$\hat{y}_{i(j-1)}^1 - \hat{y}_{ij}^1 + \hat{y}_{i(j-1)}^2 - \hat{y}_{ij}^3 + \hat{y}_{ij}^4 + \hat{y}_{ij}^5 \leq 0 \qquad \forall j = 1, \ldots, w_i - 1 \tag{14c}$$

$$\hat{y}^1_{i(w_i-1)} - \hat{y}^3_{iw_i} + \hat{y}^4_{iw_i} + \hat{y}^5_{iw_i} \leq 0 \tag{14d}$$

$$-\hat{y}^1_{i0}(p_{i1} - p_{i0}) + \hat{y}^6_{i0} + \hat{y}^7_{i0} \leq 0 \tag{14e}$$

$$\begin{aligned} &-\hat{y}^1_{ij}(p_{i(j+1)} - p_{ij}) - \hat{y}^2_{i(j-1)}(p_{ij} - p_{i(j-1)}) \\ &\quad + \hat{y}^6_{ij} + \hat{y}^7_{ij} \leq 0 \qquad \forall j = 1, \ldots, w_i - 1 \end{aligned} \tag{14f}$$

$$\sum_{j=0}^{w_i} p_{ij}\hat{y}^3_{ij} \leq z_i \tag{14g}$$

$$\sum_{j=0}^{w_i} \hat{y}^3_{ij} \leq 1 \tag{14h}$$

$$\hat{\mathbf{y}}^1_i, \hat{\mathbf{y}}^5_i, \hat{\mathbf{y}}^7_i \leq 0 \tag{14i}$$

$$\hat{\mathbf{y}}^2_i, \hat{\mathbf{y}}^3_i, \hat{\mathbf{y}}^4_i, \hat{\mathbf{y}}^6_i \geq 0. \tag{14j}$$

5.3 MILP Formulation for (3) using TEs Trained to Predict Outcomes ξ

Combining (8), (9), and (12), we can reformulate (3) as the following MIP:

$$\begin{aligned} (\mathcal{F}) \qquad \max \quad & \sum_{i \in M} \sigma_i \Psi_i(\hat{\mathbf{y}}^{1:7}_i) \\ & (\hat{\mathbf{y}}^1_i, \hat{\mathbf{y}}^2_i, \hat{\mathbf{y}}^3_i, \hat{\mathbf{y}}^4_i, \hat{\mathbf{y}}^5_i, \hat{\mathbf{y}}^6_i, \hat{\mathbf{y}}^7_i, z_i) \in \mathcal{R}_i, \quad \forall i \in M \\ & (\bar{\mathbf{x}}, \mathbf{z}, \mathbf{q}) \in Q(\mathcal{T}'). \end{aligned}$$

The expanded form of this model, after projecting out $\mathbf{z}$, can be found in Section OS.1C. Model $\mathcal{F}$ is the formulation we implement and solve in the following section for a WV case study. This study demonstrates its practical relevance and tractability.

6 Results and Discussion

In this section, we apply $\mathcal{F}$ to counties in WV. All experiments in this section were performed on an Apple M2 processor with 8.00 GB of RAM, using Python 3.13.2. The optimization model was solved using GUROBI 11.0.0. We assess the performance of our budget allocation model by implementing $\mathcal{F}$ for 12 southern WV for the years 2020, 2022 and 2023. Year 2021 is excluded from our analysis due to data inconsistencies caused by COVID-19. This region has commonly been identified as the most severely impacted by the opioid epidemic and is at the epicenter of the state's energy and economic transition - long recognized as a source of the epidemic. Before solving $\mathcal{F}$ for a (county, year) pair (c, y), we train two separate tree ensemble models, one for ORDs and the other for CIs. Based on the Arellano-Bond model results, the TE training features for predicting ORDs are previous year's $(y-1)$ ORDs, budget allocations in all six expenditure categories, *Median Household Income*, and total annual population for county c.

For predicting CIs, the TE training features include CIs from year $y-2$ and, from year $y-1$, budget allocations in all six expenditure categories, *Median Household Income*, *Median Gross Rent*, *Uninsured Population*, and total population for county c. We retain all six expenditure categories despite some showing no statistically significant causal relationship with ORDs or CIs. This is because of the budget constraint (7) in $\mathcal{F}$ that upper-bounds the total allocation across all categories. Additionally, there may also exist bounds on how much each category's allocation can deviate from its existing level.

Unlike the Arellano-Bond model, we do not apply a logarithmic transformation to the six expenditure categories. Our preliminary experiments showed no significant differences in the prescriptive results obtained with or without the transformation. This is likely because $\exp(\cdot)$ and $log(\cdot)$, being non-decreasing functions, preserve the order of values, resulting in similar splits in the trained TEs. The training data consists of the above-described information from all counties in WV for the years 2012–2023. We use scikit-learn to train TEs and apply RandomizedSearchCV to tune the hyper-parameters. Trained on expenditure and socio-economic variables with statistically significant causal links to ORDs and CIs, TEs show strong predictive performance across three metrics and consistently match or exceed Linear Regression, Support Vector Regression, and kNN, effectively capturing the budget allocation–outcome relationships, thereby justifying their use in this study. Details on tuning and predictive accuracy of TEs can be found in Section OS.1D.

In Assumption 1b, we assume that policymakers' utility is a non-decreasing function of the outcomes, meaning they favor higher outcome values. Because lower values of ORDs and CIs are desirable in our context, we transform these outcomes into normalized metrics $\tilde{\xi}_i(\mathbf{x})$ for each $i \in M$ to reflect this preference structure. Specifically, we define the normalized outcomes as:

$$\tilde{\xi}_i(\mathbf{x}) = \frac{\eta \xi_i^{max} - \xi_i(\mathbf{x})}{\xi_i^{max} - \xi_i^{min}} \qquad \forall i \in M, \tag{15}$$

where ξ_i^{max} and ξ_i^{min} represent the maximum and minimum observed values of outcome i in the dataset, respectively. The multiplier $\eta > 1$ in the numerator is included to ensure that $\tilde{\xi}_i(\mathbf{x})$ remains strictly positive, since the derivative of $\tilde{u}_i(\hat{t}, \lambda)$, as defined in (11), is undefined at $\hat{t} = 0$ for $\lambda \in [0.41, 0.68]$. In our computational experiments, we use $\eta = 1.2$. Additionally, for each outcome $i \in M$, we first normalize the leaf values in $\mathcal{T}_i$ according to (15), then sort the distinct normalized values in ascending order. These sorted values form the partition $\theta_i^1 = p_{i0} < p_{i1} < p_{i2} < \cdots < p_{iw_i} = \theta_i^2$ that is used in $\mathcal{F}$, where the total number of unique leaf values in $\mathcal{T}_i$ is $w_i + 1$.

We assess our budget allocation model, $\mathcal{F}$, by applying it to 12 counties in southern WV: Boone, Cabell, Raleigh, Fayette, Kanawha, Lincoln, Logan, McDowell, Mercer, Mingo, Wayne, and Wyoming. These counties are among the hardest hit by the opioid crisis with the highest average ORDs per 100,000 people from 2012 to 2023. A major driver of the opioid crisis in these counties is the long-term decline of coal as the foundation of their economies (Thompson et al. 2020, Young et al. 2023). In these

areas, coal is closely linked not only to residents' well-being and health, but also to the ability of local governments to provide essential services through coal severance tax revenues. In WV, these revenues are allocated to counties based on the amount of coal they produce (Hoy 2023). Substantial portions of local budgets for law enforcement and public health depend on these funds. These 12 counties lie squarely within the "coal fields" — the region most reliant on the coal industry and the tax revenues it generates. The industry's decline has left behind people with physical and mental harm, as well as underfunded government services. Together, these conditions fuel the demand for opioids while making it harder for local governments to respond effectively.

In this analysis, we evaluate three sets of coefficient values for $\{\sigma_{\texttt{ORDs}}, \sigma_{\texttt{CIs}}\}$: $\{1, 1\}$, $\{100, 1\}$, and $\{1, 100\}$, each representing a different policy emphasis. We label $\{1, 1\}$ as the *Balanced Priority Policy* (BPP), where $\mathcal{F}$ aims to achieve a solution that gives equal importance to the utility obtained from reducing both ORDs and CIs. We refer to the setting $\{100, 1\}$ where the focus is primarily on minimizing ORDs as the *Opioid Priority Policy* (OPP). Conversely, we refer to the setting $\{1, 100\}$ where reducing CIs is the higher priority as the *Crime Priority Policy* (CPP). In the RBA constraints, we set $\kappa^- = \kappa^+ = \delta\hat{x}_k$ for each $k \in N$ with $\hat{x}_k > 0$ and consider three values of $\delta = \{0.1, 0.3, 0.5\}$. A value of $\delta = 0.1$ restricts deviations from the reference budget allocations $\hat{x}_k$ for $k \in N$ to just 10%, thereby limiting the flexibility of $\mathcal{F}$ in identifying solutions that achieve lower ORDs and CIs. In contrast, setting $\delta = 0.5$ permits up to 50% deviation from the reference allocations, offering $\mathcal{F}$ significantly more flexibility in optimizing outcomes. We do not include RBA constraints for expenditure categories with $\hat{x}_k = 0$ because there are only a few such categories in our dataset. Thus, the RBA constraints for categories with $\hat{x}_k > 0$, combined with the overall budget constraint defining $\mathbf{X}$, implicitly limit the allocation to expenditure categories with $\hat{x}_k = 0$.

The remainder of this section is structured as follows. Section 6.1 analyzes the reductions in ORDs and CIs that $\mathcal{F}$ suggests are possible. Section 6.2 examines the budget allocations it recommends.

6.1 Reduction in ORDs and CIs

In this section, we compare the predicted values of ORDs and CIs (denoted as $\texttt{ORDs}_{opt}$ and $\texttt{CIs}_{opt}$) obtained from the budget allocations produced by $\mathcal{F}$ with the corresponding observed values, $\texttt{ORDs}_{obs}$ and $\texttt{CIs}_{obs}$. To quantify the improvement, we compute the percentage reductions as $\zeta^{\texttt{ORDs}} = \frac{\texttt{ORDs}_{obs} - \texttt{ORDs}_{opt}}{\texttt{ORDs}_{obs}} \times 100$ and $\zeta^{\texttt{CIs}} = \frac{\texttt{CIs}_{obs} - \texttt{CIs}_{opt}}{\texttt{CIs}_{obs}} \times 100$. We begin by focusing on a single county, Cabell. Cabell County has experienced highest per-capita rates of ORDs and high per-capita rates of CIs, along with relatively low per-capita total budget compared to other counties in WV. These conditions make Cabell County a strong candidate for applying $\mathcal{F}$.

Figure 2 plots the values of $\zeta^{\texttt{ORDs}}$ and $\zeta^{\texttt{CIs}}$ for Cabell county across the years 2020, 2022, and 2023,

under all three *priority policies* (BPP, OPP, and CPP), for each of the three δ values. As shown in Figures 2a and 2b, the values of $\zeta^{\texttt{ORDs}}$ and $\zeta^{\texttt{CIs}}$ for BPP across all years and δ levels (excluding $\delta = 0.1$ in 2022) range from 8.72% to 28.41% and from 3.79% to 23.62%, respectively. This indicates that the budget allocations produced by $\mathcal{F}$ are predicted to result in substantially lower ORDs and CIs than the observed values. Given the strong predictive accuracy of TEs and the causal relationships identified in Section 4, these findings underscore the potential of model $\mathcal{F}$ to aid policymakers in adjusting budget allocation decisions to reduce ORDs and CIs.

Furthermore, as seen in Figures 2c and 2f, the values of $\zeta_{\texttt{ORDs}}$ and $\zeta_{\texttt{CIs}}$ under OPP and CPP, respectively, are non-decreasing as δ increases. This trend is expected since higher δ values allow $\mathcal{F}$ greater flexibility in adjusting budget allocations across expenditure categories, enabling the model to identify allocations that better optimize the targeted outcome. When comparing across *priority policies*, we observe that $\zeta_{\texttt{ORDs}}$ improves under OPP relative to BPP, while $\zeta_{\texttt{CIs}}$ declines. For example, in 2022 with $\delta = 0.3$, $\zeta_{\texttt{ORDs}}$ rises from 23.5% under BPP to 25.3% under OPP, whereas $\zeta_{\texttt{CIs}}$ falls from 3.8% to –5.6%. A similar but reversed pattern emerges when comparing CPP with BPP. Under CPP, $\zeta_{\texttt{CIs}}$ increases, while $\zeta_{\texttt{ORDs}}$ declines. Specifically, for 2022 and $\delta = 0.3$, $\zeta_{\texttt{CIs}}$ increases from 3.8% under BPP to 7.7% under CPP, while $\zeta_{\texttt{ORDs}}$ drops from 23.5% to 8.9%. These results suggest that, for Cabell County, a trade-off may exist between reducing ORDs and CIs: improvements in one outcome may come at the cost of the other.

Such trade-offs in outcomes are well-known in the study of political science and public policy (Breunig and Busemeyer 2012). In the counties we examine, high rates of co-morbidities have increased the likelihood that public policies across different substantive areas become closely interconnected. In Cabell County, for instance, millions of opioid pills were shipped to Huntington, WV, and the surrounding areas (Voelker 2018). During this same period, the county experienced significant levels of co-morbidities, particularly obesity and cigarette smoking, making it even more difficult to balance public health needs with drug interdiction efforts. These health challenges constrain policy flexibility, tightening the linkage between competing priorities.

Next, we examine whether the above insights also apply to the other 11 counties of southern WV. We found that all 11 other counties exhibit trends similar to those observed in Cabell County. In most cases, $\mathcal{F}$ identifies budget allocations where the predicted ORDs and CIs are lower than the observed values. Moreover, as δ increases, $\zeta^{\texttt{ORDs}}$ and $\zeta^{\texttt{CIs}}$ tend to improve under the OPP and CPP, respectively, whereas under the BPP policy, at least one of the two metrics improves. Across various *priority policies*, similar trade-offs to those seen in Cabell County emerge. In most of the other 11 counties analyzed, $\zeta^{\texttt{ORDs}}$ tends to be higher under OPP than under BPP, whereas $\zeta^{\texttt{CIs}}$ is generally higher under CPP than BPP. This suggests that prioritizing improvements in one outcome often would come at the cost of reducing gains in the other, a pattern we found to be common throughout southern WV. These trade-offs are well-documented

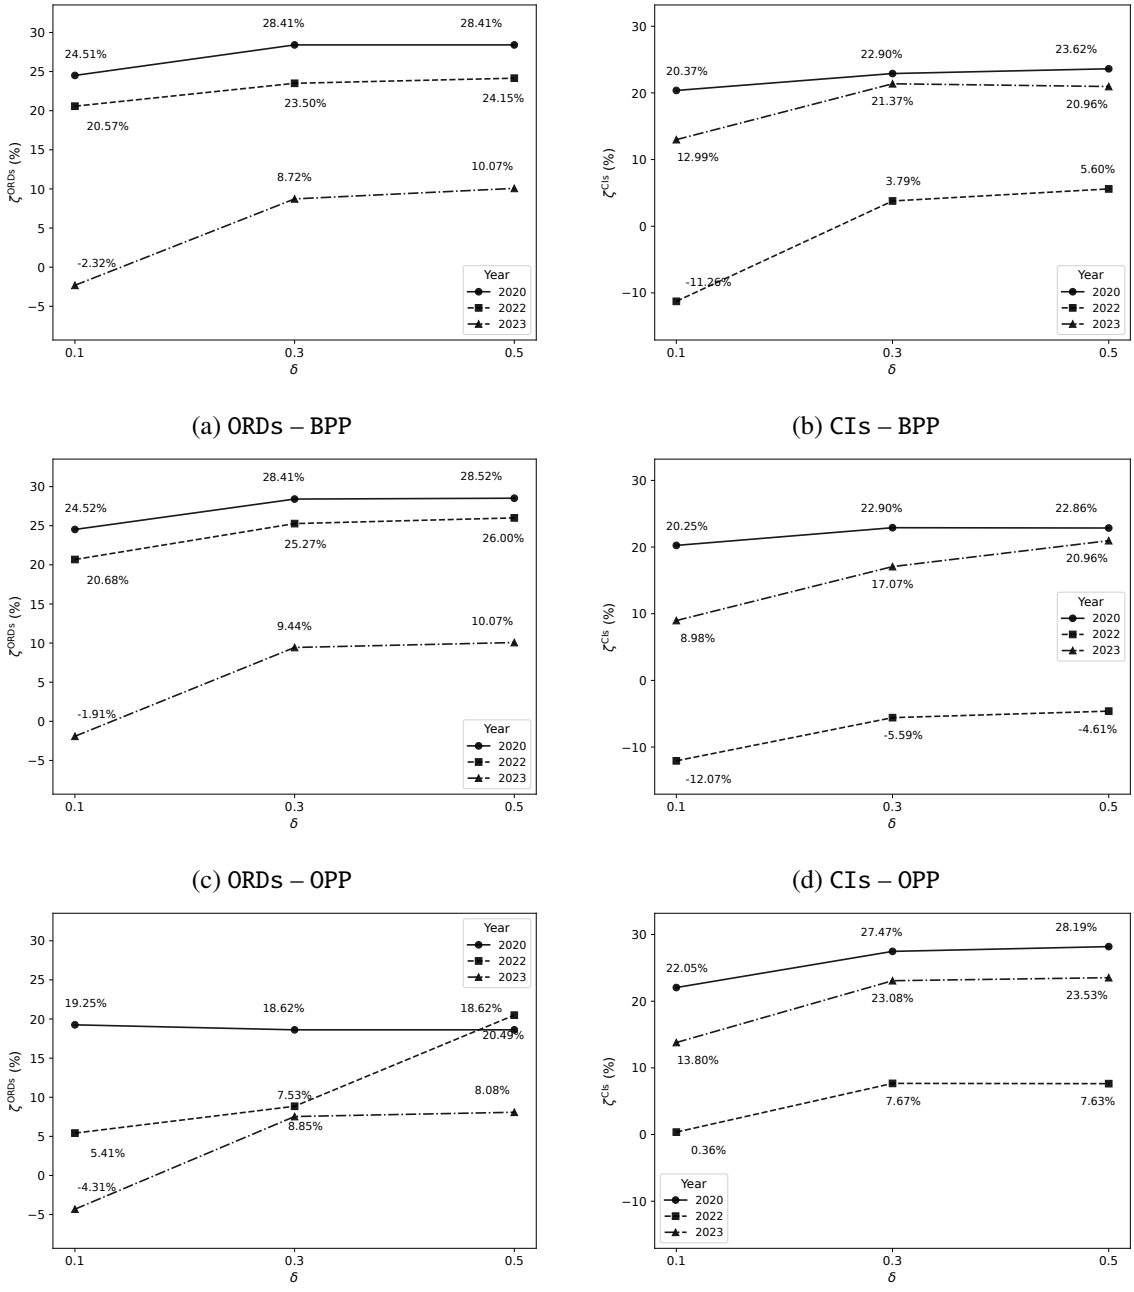

(a) ORDs – BPP

(b) CIs – BPP

(c) ORDs – OPP

(d) CIs – OPP

(e) ORDs – CPP

(f) CIs – CPP

Figure 2: Percentage reductions in ORDs and CIs ($\zeta^{\texttt{ORDs}}$ and $\zeta^{\texttt{CIs}}$) across various δ values under *priority policies* for Cabell County.

in theoretical decision models of governance and public policy and are supported by empirical evidence (Breunig and Busemeyer 2012). We refer to Figures OS.3 and OS.4 for graphs summarizing our findings for all 12 counties.

6.2 Changes in Budget Allocations

We now examine the budget allocations that lead to the predicted ORDs and CIs values described in Section 6.1. For each budget category $k \in N$, an optimal solution to $\mathcal{F}$ provides a budget range from the set $\bar{L}_k$. Therefore, instead of a single allocation, $\mathcal{F}$ produces an interval for each category. This offers policymakers greater flexibility in selecting budget allocations, subject to the overall budget constraint. We denote this interval for category k, county c and year y as $(\mathtt{LB}^y_{ck}, \mathtt{UB}^y_{ck})$.

Let $\tilde{x}^y_{ck}$ represent the observed budget allocation for county c in category k for year y. We assess the recommended budget adjustments through the interval $(\tilde{\Delta}^y_{ck}, \hat{\Delta}^y_{ck})$ where $\tilde{\Delta}^y_{ck} = \frac{\mathtt{LB}^y_{ck} - \tilde{x}^y_{ck}}{b^y_c} \times 100$ and $\hat{\Delta}^y_{ck} = \frac{\mathtt{UB}^y_{ck} - \tilde{x}^y_{ck}}{b^y_c} \times 100$. Here, b^y_c denotes the total budget of county c in year y, whereas $\tilde{\Delta}^y_{ck}$ and $\hat{\Delta}^y_{ck}$ reflects the minimum and maximum percentage change, respectively, in the allocation to category k relative to the county's total budget for that year.

Table 2 reports the values of $(\tilde{\Delta}^y_{ck}, \hat{\Delta}^y_{ck})$ for Cabell County in 2023 under all combinations of δ levels and *priority policies*. The last two rows of the table display the values of $\tilde{x}^y_{ck}$ and the total county budget b^y_c for 2023. The table's final column presents the value of $\kappa^y_c = \frac{\sum_{k \in N} \mathtt{UB}^y_{ck}}{b^y_c}$ for each δ level and *priority policy*, representing the maximum fraction of the total budget that can be allocated based on the bounds prescribed by $\mathcal{F}$. As shown in the table, $\mathcal{F}$ recommends significant changes in three expenditure categories: CR, GG, and PS. Under the BPP policy, as δ increases from 0.1 to 0.3, $\mathcal{F}$ suggests a substantial decrease in allocations to CR and GG, and an increase in spending on PS. These adjustments lead to an improvement in outcomes, with $\zeta^{\mathtt{ORDs}}$ rising from -2.32% to 8.72% and $\zeta^{\mathtt{CIs}}$ from 12.99% to 21.37%; see Figures 2a and 2b. These findings align with the Arellano-Bond model results in Table 1, where CR has a statistically significant positive coefficient for ORDs, and PS has a statistically significant negative coefficient for CIs. This suggests that reducing spending on CR and increasing investment in PS may help reduce ORDs and CIs, respectively. Although GG shows a statistically significant negative relationship with ORDs in the Arellano-Bond model, $\mathcal{F}$ still recommends decreasing its allocation. This is due to the budget constraint, which the Arellano-Bond model does not account for. Moreover, when δ increases to 0.5, a shift occurs: $\mathcal{F}$ recommends reducing funding to PS and increasing it for GG due to the budget limitation. This change results in a further increase in $\zeta^{\mathtt{ORDs}}$ to 10.07%, but a slight decrease in $\zeta^{\mathtt{CIs}}$ to 20.96% compared to the outcomes at $\delta = 0.3$.

Under the OPP policy, which prioritizes reducing ORDs, $\mathcal{F}$ recommends decreasing allocations to both PS and CR, a recommendation that aligns with the Arellano-Bond model results in Table 1. However,

an interesting pattern emerges with respect to GG as δ increases. Specifically, $\mathcal{F}$ suggests a reduction in GG spending when $\delta = 0.3$, followed by an increase when $\delta = 0.5$. As illustrated in Figure 2c, the values of $\zeta^{\texttt{ORDs}}$ for $\delta = 0.3$ and 0.5 are quite close—9.44% and 10.07%, respectively. To attain these reductions in ORDs, $\mathcal{F}$ identifies two distinct strategies: one involves simultaneously making a substantial cut in CR spending while increasing allocation to GG, and the other reflects a scenario where a significant reduction in CR is not feasible. The first strategy becomes viable only when a higher flexibility level (*i.e.*, $\delta = 0.5$) allows for greater reallocation across categories. When δ is smaller (*e.g.*, 0.3), the scope for cutting CR is limited, prompting $\mathcal{F}$ to reduce spending on GG instead. This illustrates how $\mathcal{F}$ accounts for interdependencies between budget categories—recognizing that achieving notable improvements in outcomes may require coordinated, simultaneous adjustments across multiple expenditure areas.

This seemingly counterintuitive result can be attributed to the way public institutions process information—through complex, nonlinear mechanisms that involve threshold effects. Governments often respond only after a problem becomes severe, at which point their actions may be disproportionate. This indicates that the government action only occurs after an outcome crosses a certain threshold, and such thresholds are not always triggered in a gradual, incremental fashion. As a result, influencing outcomes often requires substantial budget cuts or expansions (Workman et al. 2009, Workman and Thomas 2025).

Under the CPP policy, which emphasizes reducing CIs, increasing δ from 0.1 to 0.5 results in greater budget allocation to PS and reduced spending on CR and GG. This reallocation leads to an increase in $\zeta^{\texttt{CIs}}$ from 13.80% to 23.53%; see Figure 2f. These patterns are in line with the Arellano-Bond model results in Table 1, where PS has a statistically significant negative coefficient with respect to CIs, suggesting that increased spending in this category is associated with reduced crime incidents. This also aligns with intuitive expectations—boosting investment in public safety is a reasonable approach when the goal is to lower crime.

Next, we extend our analysis to include the 11 other counties in southern WV. Our goal is to examine whether our model reveals county-wise differences in budget allocations across southern WV. We found that most counties direct the most of their funds toward GG and PS, whereas CP receives only a minimal portion; see Table OS.4. This spending pattern stands in contrast to the Arellano-Bond model findings (Table 1), which identify CP as having a statistically significant negative association with both ORDs and CIs. In line with this evidence, the most substantial budget reallocation recommendations made by $\mathcal{F}$ across the 12 counties focus on CP, GG, and PS. In many of the 12 counties, $\mathcal{F}$ suggests notable increases in funding for CP. Given the constraint of total budget, the increased funding for CP is primarily achieved by reducing allocations to GG. This is mainly due to two reasons: first, GG currently receives the highest share of the total budget, making it the most practical source for reallocation; and second, GG does not have a statistically significant effect on CIs (see Table 1). Consequently, under the CPP policy, the model

Priority Policy	δ	CP	CR	GG	HS	PS	SS	κ_c^y
BPP	0.10	(-0.021, 0.021)	(-1.077, -0.964)	(-5.502, -4.580)	(0.009, 0.020)	(3.192, 3.381)	(0.034, 0.036)	0.98
	0.30	(-0.062, -0.030)	(-3.138, -3.110)	(-5.574, -4.580)	(0.009, 0.030)	(5.494, 6.379)	(0.259, 0.263)	0.99
	0.50	(-0.103, -0.030)	(-5.052, -5.048)	(8.588, 10.616)	(-0.026, -0.024)	(-3.998, -3.876)	(0.263, 0.282)	1.02
OPP	0.10	(-0.021, 0.021)	(-1.077, -0.964)	(3.188, 3.505)	(0.009, 0.020)	(-3.381, -3.375)	(0.142, 0.144)	0.99
	0.30	(-0.062, -0.030)	(-2.092, -1.999)	(-4.131, -4.108)	(-0.026, -0.024)	(-3.998, -3.876)	(0.263, 0.282)	0.90
	0.50	(-0.103, -0.030)	(-5.052, -5.048)	(8.588, 10.616)	(-0.027, -0.026)	(-3.998, -3.876)	(0.263, 0.282)	1.02
CPP	0.10	(-0.021, 0.021)	(-0.791, 0.359)	(-5.502, -4.580)	(0.009, 0.020)	(-2.983, -2.600)	(0.651, 0.910)	0.94
	0.30	(-0.062, -0.030)	(-3.138, -3.110)	(-5.574, -4.580)	(0.009, 0.030)	(4.671, 5.143)	(0.651, 0.910)	0.98
	0.50	(-0.103, -0.030)	(-5.048, -5.033)	(-4.131, -4.108)	(0.088, 0.098)	(5.143, 5.156)	(1.522, 1.526)	0.98
$\tilde{x}_{ck}^y$		\$60,079	\$3,154,678	\$16,114,350	\$57,147	\$9,902,121	\$0	
b_c^y		\$29,288,375						

Table 2: Minimum & maximum percentage change values ($\tilde{\Delta}_{ck}^y, \hat{\Delta}_{ck}^y$) for c = Cabell County and $y = 2023$ across δ values and *priority policies*.

recommends sharper reductions in GG funding at higher values of δ. As for PS, $\mathcal{F}$ generally advises reducing allocations across most counties, though some exceptions exist. For example, with $\delta = 0.5$, 2 counties receive an increase in PS funding under BPP, while 4 each receive an increase under OPP and CPP. In the remaining counties, $\mathcal{F}$ recommends reducing PS allocations or keeping them at existing levels. Although these reductions are generally smaller than those seen for GG, they are justified by PS receiving the second-largest share of the budget. The increases can be explained by the fact that PS is found to have a statistically significant negative impact on CIs as per the results of the Arellano-Bond model in Table 1. For graphical representation of these results, see Figures OS.5–OS.10.

These results align with the realities of local government operations and spending priorities. GG represents the pool of expenditures for running the government and receives the largest share of resources. In addressing both drug treatment and crime prevention, a persistent and shifting tension exists regarding how funds are distributed across policy domains. The model's guidance for GG points toward a redefinition of local government's central functions, indicating a major reordering of policy priorities.

Reallocating resources as suggested by $\mathcal{F}$ could strengthen the response to the drug crisis but might come at the expense of infrastructure and other essential services. Reducing PS funding often faces opposition from powerful law enforcement interests and, as our analysis shows, could result in higher crime rates. County-specific context is critical for these decisions: in areas where geographic, infrastructural, and demographic factors make supply interdiction crucial, cuts to PS could be minimal; in areas where such interdiction is unlikely or impractical, a more balanced spending strategy could be prioritized. Ultimately, counties assess both logistical and political feasibility, as effective policies must fall within an acceptable "zone of acquiescence" on these dimensions; otherwise, adoption will be slow

or absent.

7 Conclusion

In public policy broadly—and in public health and criminal justice in particular—policies are often conceptualized as one-time interventions. There are valid reasons for this: evidence-based research identifies certain policies that, if implemented, tend to improve outcomes in addressing Substance Use Disorder (SUD). Methodological conventions also reinforce this view, as many statistical frameworks are built around the assumption that a single policy is enacted at a given point in time, followed by measurable outcomes later. However, this "single-intervention" perspective does not reflect the actual decision environment of most policymakers. They must juggle numerous competing issues, each demanding attention and policy action. In practice, policymakers rely on a combination of decisions to nudge governance toward their preferred outcomes, working within the limits of their bounded rationality. From this perspective, policies aimed at any particular issue, such as SUD, are better understood as a mix of actions, rather than a single intervention. Moreover, while the intervention-based model may be grounded in science, it often overlooks the socio-economic and political realities of decision-making. Some policy choices are simply more acceptable to local officials than others. By framing public policy as a mix of tools or levers for influencing outcomes, we provide a more adaptable framework that accommodates diverse political and social contexts. This flexibility improves the likelihood of policy adoption in real-world governance.

In this paper, we develop a county-level decision support framework that integrates causal inference, machine learning, and utility optimization to inform budget allocation strategies targeting opioid-related outcomes. The proposed framework explicitly captures the complex relationships between county-level budget allocations and two key outcomes (ORDs and CIs) while accounting for the unknown or unspecified priorities that policymakers may have regarding these outcomes. We leverage causal inferences from the Arellano-Bond model to guide the selection of expenditure categories and socio-economic variables used as input features in training the TEs that predict ORDs and CIs. Subsequently, for a given county, we embed trained TEs into an MILP model that generates budget allocations across expenditure categories to optimize the policymakers' worst-case risk-averse utility. This approach enables the development of a tailored, county-level decision support tool to guide budget allocations, helping to more effectively manage the opioid crisis.

The application of our approach to the state of WV yields several important findings. The Arellano-Bond model establishes statistically significant causal relationships between budget allocations in different expenditure categories, socio-economic factors, and both ORDs and CIs. In our computational experiments, we apply $\mathcal{F}$ to 12 southern counties in WV—the area most acutely affected by the opioid

crisis—using six expenditure categories: CP, CR, GG, HS, PS, and SS, and targeting two opioid-related outcomes: ORDs and CIs. These outcomes are by far the most important for addressing the opioid crisis, representing the direct effect on individuals and families, as well as the externalities (e.g., crime). Our computational results indicate that $\mathcal{F}$ can effectively recommend budget allocations across the six expenditure categories to substantially reduce both ORDs and CIs in all 12 southern WV counties. We further evaluate the performance of $\mathcal{F}$ under three distinct *priority policies*: BPP (equal emphasis on ORDs and CIs), OPP (greater emphasis on ORDs), and CPP (greater emphasis on CIs). Our analysis reveals a trade-off between the two outcomes—greater improvements in one may lead to lower improvements in the other.

Building on the observed outcome improvements, we analyze the underlying budget allocations that drive them. Instead of providing fixed allocation values, $\mathcal{F}$ generates intervals for each expenditure category, offering policymakers flexibility in implementation. We assess these allocations under three levels of $\delta = 0.1$, 0.3, and 0.5, where δ represents the maximum allowable increase or decrease in any category relative to its current allocation. At higher δ values, the budget shifts prescribed by $\mathcal{F}$ closely mirror the relationships identified by the Arellano-Bond model. Further, we observe that $\mathcal{F}$ captures interdependencies among expenditure categories—indicating that meaningful improvements in outcomes often require coordinated, multi-category budget adjustments.

We identify three directions of future research that may improve the generalizability of the proposed framework. First, incorporating additional outcome metrics, such as the share of individuals in recovery, public confidence in law enforcement, and economic growth, can offer a more holistic view of how budget allocations help in managing the opioid crisis. Second, as each of the six expenditure categories contains multiple subcategories, the models in this paper can also guide budget allocation at a more granular level, enabling more targeted and effective micro-level decision-making. The results generated by our model can serve as the foundation for a two-step subcategory-level analysis. In the first step, the model identifies which major budget categories should receive increased or decreased funding—and by what amount—to help minimize ORDs and CIs. In the second step, the specific subcategories within each major category can have their allocations adjusted in a manner consistent with the overall changes identified in step one, while still adhering to the objective of reducing ORDs and CIs. Third, the current approach assumes policymakers are risk-averse in their decision-making. However, decision-makers may sometime demonstrate more nuanced, potentially risk-seeking preferences. A budget allocation model that relaxes this assumption could prove useful at capturing policymakers' complexities.

Data Availability Statement: The data that support the findings of this study are available from the corresponding author, upon reasonable request.

References

Anderson, R., J. Huchette, W. Ma, C. Tjandraatmadja, and J. P. Vielma (2020). Strong mixed-integer programming formulations for trained neural networks. *Mathematical Programming 183*(1), 3–39.

Ansari, S., S. Enayati, R. Akhavan-Tabatabaei, and J. M. Kapp (2024). Curbing the opioid crisis: Optimal dynamic policies for preventive and mitigating interventions. *Decision Analysis 21*(3), 165–193.

Arellano, M. and S. Bond (1991). Some tests of specification for panel data: Monte carlo evidence and an application to employment equations. *The Review of Economic Studies 58*(2), 277–297.

Baardman, L., M. C. Cohen, K. Panchamgam, G. Perakis, and D. Segev (2019). Scheduling promotion vehicles to boost profits. *Management Science 65*(1), 50–70.

Battista, N. A., L. B. Pearcy, and W. C. Strickland (2019). Modeling the prescription opioid epidemic. *Bulletin of Mathematical Biology 81*, 2258–2289.

Baucum, M., M. Harris, L. Kessler, and G. Lu (2025). Reducing overdose deaths and mitigating county disparities: Optimal allocation of substance use treatment centers. *Manufacturing & Service Operations Management 27*(3), 736–756.

Bertsimas, D., A. O'Hair, S. Relyea, and J. Silberholz (2016). An analytics approach to designing combination chemotherapy regimens for cancer. *Management Science 62*(5), 1511–1531.

Bowden, N., R. Merino, S. Katamneni, and A. Coustasse (2018). The cost of the opioid epidemic in West Virginia. Paper presented at the 54th Annual MBAA Conference, Chicago, IL.

Breunig, C. and M. R. Busemeyer (2012). Fiscal austerity and the trade-off between public investment and social spending. *Journal of European Public Policy 19*(6), 921–938.

Cesare, N. et al. (2024). Development and validation of a community-level social determinants of health index for drug overdose deaths in the healing communities study. *Journal of Substance Use and Addiction Treatment 157*, 209186.

Chen, Q. et al. (2019). Prevention of prescription opioid misuse and projected overdose deaths in the united states. *JAMA Network Open 2*(2), e187621–e187621.

Compton, W. M., C. M. Jones, and G. T. Baldwin (2016). Relationship between nonmedical prescription opioid use and heroin use. *The New England Journal of Medicine 374*, 154–163.

Dave, D., M. Deza, and B. Horn (2021). Prescription drug monitoring programs, opioid abuse, and crime. *Southern Economic Journal 87*(3), 808–848.

Deepa, C., K. Sathiya Kumari, and V. Pream Sudha (2010). A tree based model for high performance concrete mix design. *International Journal of Engineering Science and Technology 2*(9), 4640–4646.

Drug Enforcement Administration (2025). 2025 national drug threat assessment summary. Updated. US Department of Justice. DEA-DCT-DIR-007-25.

Fernández-Delgado, M., E. Cernadas, S. Barro, and D. Amorim (2014). Do we need hundreds of classifiers to solve real world classification problems? *Journal of Machine Learning Research 15*(1), 3133–3181.

Ferreira, K. J., B. H. A. Lee, and D. Simchi-Levi (2016). Analytics for an online retailer: Demand forecasting and price optimization. *Manufacturing & Service Operations Management 18*(1), 69–88.

Fischetti, M. and J. Jo (2018). Deep neural networks and mixed integer linear optimization. *Constraints 23*(3), 296–309.

Gong, G., S. G. Phillips, C. Hudson, D. Curti, and B. U. Philips (2019). Higher US rural mortality rates linked to socioeconomic status, physician shortages, and lack of health insurance. *Health Affairs 38*(12), 2003–2010.

Hastie, T., R. Tibshirani, and J. Friedman (2009). *The Elements of Statistical Learning*. Springer.

Hedegaard, H., M. Warner, and A. M. Minino (2017). Drug overdose deaths in the United States, 1999–2016. *NCHS Data Brief 294*, 1–8.

Herrera, M., L. Torgo, J. Izquierdo, and R. Pérez-García (2010). Predictive models for forecasting hourly urban water demand. *Journal of Hydrology 387*(1-2), 141–150.

Holland, K., L. DePadilla, D. Gervin, E. Parker, and M. Wright (2021). The evolution of the overdose epidemic and CDC's research response: a commentary. *Drugs Context 10*(2021-8-2).

Holt, C. A. and S. K. Laury (2002). Risk aversion and incentive effects. *American Economic Review 92*(5), 1644–1655.

Hoy, K. A. (2023). Asymmetric responses to severance tax changes: Coal production in West Virginia. *Energy Economics 125*, 106840.

Hu, J. and S. Mehrotra (2015). Robust decision making over a set of random targets or risk-averse utilities with an application to portfolio optimization. *IIE Transactions 47*(4), 358–372.

İskenderoğlu et al. (2020). Comparison of support vector regression and random forest algorithms for estimating the sofc output voltage by considering hydrogen flow rates. *International Journal of Hydrogen Energy 45*(60), 35023–35038.

James, G., D. Witten, T. Hastie, R. Tibshirani, et al. (2013). *An introduction to statistical learning*, Volume 112. Springer.

Jones, B. D. (2003). Bounded rationality and political science: Lessons from public administration and public policy. *Journal of Public Administration Research and Theory 13*(4), 395–412.

Jones, B. D. and F. R. Baumgartner (2005). *The Politics of Attention: How Government Prioritizes Problems*. University of Chicago Press.

Jones, M. R., O. Viswanath, J. Peck, A. D. Kaye, J. S. Gill, and T. T. Simopoulos (2018). A brief history of the opioid epidemic and strategies for pain medicine. *Pain and Therapy 7*, 13–21.

Kahneman, D. and A. Tversky (1979). Prospect theory: An analysis of decision under risk. *Econometrica 47*(2), 263–291.

Kim, J., J.-P. P. Richard, and M. Tawarmalani (2025). A reciprocity between tree ensemble optimization and multilinear optimization. *Operations Research 73*(5), 2610–2626.

Kolodny, A. et al. (2015). The prescription opioid and heroin crisis: a public health approach to an epidemic of addiction. *Annual Review of Public Health 36*, 559–574.

Lindenfeld, Z., D. Silver, A. I. Mauri, and M. W. Rothbart (2025). Beyond social determinants: Fiscal determinants of overdose death in US counties, 2017–2020. *Social Science & Medicine 364*, 117529.

Luo, J. and B. Stellato (2024). Frontiers in operations: Equitable data-driven facility location and resource allocation to fight the opioid epidemic. *Manufacturing & Service Operations Management 26*(4), 1229–1244.

Mars, S. G., P. Bourgois, G. Karandinos, F. Montero, and Ciccarone (2014). Every 'never' i ever said came true': transitions from opioid pills to heroin injecting. *International Journal of Drug Policy 25*, 257–66.

Marshall, B., M. K. Bland, R. Hulla, and R. J. Gatchel (2019). Considerations in addressing the opioid epidemic and chronic pain within the USA. *Pain Management 9*, 131–138.

Mas-Colell, A., M. D. Whinston, J. R. Green, et al. (1995). *Microeconomic theory*, Volume 1. OUP, New York.

Mattson, C. L., L. J. Tanz, K. Quinn, M. Kariisa, P. Patel, and N. L. Davis (2021). Trends and geographic patterns in drug and synthetic opioid overdose deaths–United States, 2013–2019. *Morbidity and Mortality Weekly Report 70*(6), 202.

Merino, R., N. Bowden, S. Katamneni, and A. Coustasse (2019). The opioid epidemic in West Virginia. *Health Care Management 38*, 187–195.

Mišić, V. V. (2020). Optimization of tree ensembles. *Operations Research 68*(5), 1605–1624.

Mistry, M., D. Letsios, G. Krennrich, R. M. Lee, and R. Misener (2021). Mixed-integer convex nonlinear optimization with gradient-boosted trees embedded. *INFORMS Journal on Computing 33*(3), 1103–1119.

Office of National Drug Control Policy (2022). National drug control budget: Federal year 2022 highlights. `https://bidenwhitehouse.archives.gov/wp-content/uploads/2023/03/NDCS-FY2022-Budget-Summary.pdf`.

Ogunade, A. C., M. Kruger, D. Eades, C. Fitch, and M. Boki (2022). Assessment of west virginia public issues: Perceptions of key decisionmakers. Report, West Virginia University Extension.

Pitt, A. L., K. Humphreys, and M. L. Brandeau (2018). Modeling health benefits and harms of public policy responses to the US opioid epidemic. *American Journal of Public Health 108*(10), 1394–1400.

Rao, I. J., K. Humphreys, and M. L. Brandeau (2021). Effectiveness of policies for addressing the us opioid epidemic: a model-based analysis from the Stanford-Lancet commission on the north american opioid crisis. *The Lancet Regional Health–Americas 3*.

Sim, Y. (2023). The effect of opioids on crime: Evidence from the introduction of oxycontin. *International Review of Law and Economics 74*, 106136.

Sistani, F., M. R. de Bittner, and F. T. Shaya (2023). Social determinants of health, substance use, and drug overdose prevention. *Journal of the American Pharmacists Association 63*(2), 628–632.

StataCorp (2025). *Stata Statistical Software: Release 19*. College Station, TX: StataCorp LLC.

Thompson, J. R., S. L. Creasy, C. F. Mair, and J. G. Burke (2020). Drivers of opioid use in Appalachian Pennsylvania: Cross-cutting social and community-level factors. *International Journal of Drug Policy 78*, 102706.

U.S. Bureau of Economic Analysis (2025). Gross domestic product: Implicit price deflator (gdpdef). FRED, Federal Reserve Bank of St. Louis.

U.S. Census Bureau (2025). Data.census.gov. Accessed: 2025-08-13.

Voelker, R. (2018). A day in the life: facing the opioid epidemic in Huntington, West Virginia. *JAMA 319*(14), 1423–1424.

West Virginia State Auditor (2025). https://www.wvsao.gov/LocalGovernment/Default.

White, E. and C. Comiskey (2007). Heroin epidemics, treatment and ode modelling. *Mathematical Biosciences 208*(1), 312–324.

Wooldridge, J. M. (2010). *Econometric analysis of cross section and panel data*. MIT press.

Workman, S. (2015). *The Dynamics of Bureaucracy in the U.S. Government*. Cambridge University Press.

Workman, S., B. D. Jones, and A. E. Jochim (2009). Information processing and policy dynamics. *Policy Studies Journal 37*(1), 75–92.

Workman, S. and H. F. Thomas (2025). The other side of chaos COVID-19, federal spending, and local government volatility. *Journal of European Public Policy*.

Workman, S., H. F. Thomas, and C. N. Connor (2025). Federalism, punctuated equilibrium, and local budget dynamics: Findings from Appalachia. Available at Institute for Policy Research and Public Affairs.

Yaemsiri, S., J. M. Alfier, E. Moy, L. M. Rossen, B. Bastian, J. Bolin, A. O. Ferdinand, T. Callaghan, and M. Heron (2019). Healthy people 2020: rural areas lag in achieving targets for major causes of death. *Health Affairs 38*(12), 2027–2031.

Young, T., J. Baka, Z. He, S. Bhattacharyya, and Z. Lei (2023). Mining, loss, and despair: exploring energy transitions and opioid use in an Appalachian coal community. *Energy Research & Social Science 99*, 103046.

Online Supplement

OS.1: Supplementary Material

OS.1A Graphical Description of Collected Data

In this section, we use maps of West Virginia to visualize geographical patterns in our data on county budgets and opioid crisis outcomes. Figures OS.1a to OS.1d display the average population, `ORDs` per 100,000 residents, `CIs` per 100,000 residents, and average budget per 100,000 residents, respectively, for each county in `WV` over the period 2012–2023. As shown in Figure OS.1b, the southern counties, such as McDowell, Logan, Wyoming, and Cabell, exhibit the highest average opioid mortality rates, pointing to stark regional disparities. Similarly, Figure OS.1c indicates that southern counties such as Kanawha, Raleigh, Fayette, and Cabell report the highest average crime incident rates per 100,000 residents. Finally, Figure OS.1d shows that counties bearing the highest burden of `ORDs` and `CIs`, such as McDowell, Cabell, and Kanawha, have some of the lowest per capita budgets. The scarcity of monetary resources in many counties makes clear the importance of allocating county budgets efficiently to effectively combat the opioid crisis.

OS.1B Background on *Tree Ensembles* (`TEs`)

To inform budget allocations, our model must identify from data the relationships between allocations and outcomes. While various machine learning models could be applied, in this paper we employ `TEs` because they have proven successful in many applications (Ferreira et al. 2016, Deepa et al. 2010, Herrera et al. 2010) and are among the best predictive models across various datasets (Fernández-Delgado et al. 2014).

`TEs` are machine learning models that are used for classification (when the output takes values in a finite discrete set) and regression (when the output variable takes continuous values). `TEs` are a collection of a decision trees where each decision tree can be viewed as a function that, given any value for the independent variables $(X_1, \ldots, X_n)$, produces the value of a dependent variable Y by answering a cascade of queries, each checking whether a single input variable falls below or exceeds a specified threshold. Because of the cascading nature of the queries, this function can be visualized as a tree. The output of a decision tree is obtained by following a sequence of queries from the root to a leaf node and returning the value associated with that leaf during training. A decision tree models the dependent variable Y as a piecewise constant function over the input space. Each leaf of the tree corresponds to a hyper-rectangular region defined by the threshold-based conditions on the path leading to that leaf. The output of a `TE` is computed by combining the individual predictions of multiple decision trees, typically through averaging. `TEs` are preferred over single decision trees as they exhibit much-enhanced prediction accuracy.

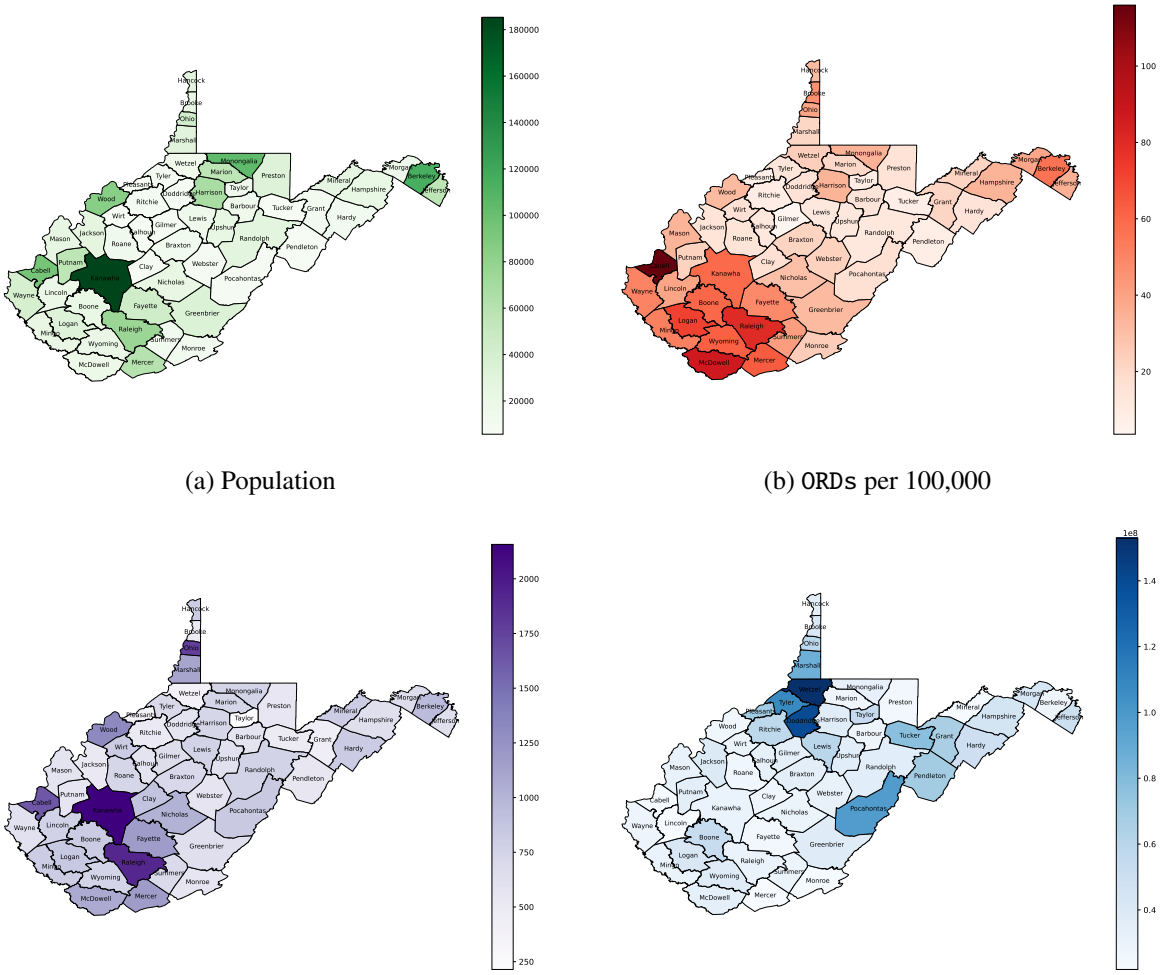

(a) Population

(b) ORDs per 100,000

(c) CIs per 100,000

(d) Budget per 100,000

Figure OS.1: Population and average per-capita metrics across WV counties from 2012–2023. Subplots show: (a) population, (b) ORDs per 100,000, (c) CIs per 100,000, and (d) budget per 100,000.

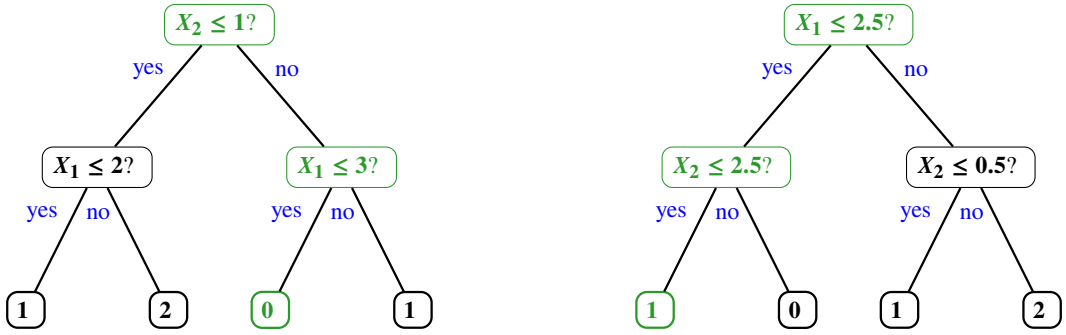

Figure OS.2: Tree ensemble predicting value $\frac{1}{2}$ (average of green leaves' values) when $(X_1, X_2) = (1.2, 2.1)$.

To give a concrete example, we represent in Figure OS.2, a TE with two trees that predicts the value of a dependent (output) variable Y from independent (input) variables X_1 and X_2. In the context of our application, one may think of X_1 as the amount of money spent on PS, X_2 as the amount money spent on

HS, and Y as the number of opioid deaths. When $(X_1, X_2) = (1.2, 2.1)$, the trees of the TE of Figure OS.2 predict values of Y equal to 0 and 1, respectively. To generate these predictions, each tree is traversed from the root to a leaf node by selecting, at each split, the branch corresponding to the condition satisfied by the input $(X_1, X_2) = (1.2, 2.1)$. Generally, a TE outputs the average of predictions from the individual trees, *i.e.*, $\frac{1}{2}$ in our example. In practice, TEs typically consist of numerous trees that are deeper and depend on a greater number of variables than those illustrated in Figure OS.2.

OS.1C Extensive Form of $\mathcal{F}$

The budget allocation model $\mathcal{F}$ we proposed in Section 5.3 is obtained by reformulating (3). This reformulation integrates the mixed-integer set from Section 5.1, which encodes how tree ensembles capture the relationship between budget allocations and outcomes, with the linear constraints from Section 5.2, which characterize the worst-case utility maximization problem. Next, we present the complete set of constraints for $\mathcal{F}$ in a single display.

$$
\begin{aligned}
(\mathcal{F}) \quad \max \; & \sum_{i \in M} \sum_{j=0}^{w_i} \sigma_i \begin{pmatrix} \tilde{u}_i(p_{ij}, 0.41)\hat{y}^4_{ij} + \tilde{u}_i(p_{ij}, 0.68)\hat{y}^5_{ij} \\ +\tilde{u}'_i(p_{ij}, 0.41)\hat{y}^6_{ij} + \tilde{u}'_i(p_{ij}, 0.68)\hat{y}^7_{ij} \end{pmatrix} & \\
& -\hat{y}^1_{i0} - \hat{y}^2_{i0} - \hat{y}^3_{i0} + \hat{y}^4_{i0} + \hat{y}^5_{i0} \leq 0 && \forall i \in M \\
& \hat{y}^1_{i(j-1)} - \hat{y}^1_{ij} + \hat{y}^2_{i(j-1)} - \hat{y}^2_{ij} - \hat{y}^3_{ij} + \hat{y}^4_{ij} + \hat{y}^5_{ij} \leq 0 && \forall i \in M, \forall j = 1, \ldots, w_i - 2 \\
& \hat{y}^1_{i(j-1)} - \hat{y}^1_{ij} + \hat{y}^2_{i(j-1)} - \hat{y}^3_{ij} + \hat{y}^4_{ij} + \hat{y}^5_{ij} \leq 0 && \forall i \in M, \forall j = 1, \ldots, w_i - 1 \\
& \hat{y}^1_{i(w_i-1)} - \hat{y}^3_{iw_i} + \hat{y}^4_{iw_i} + \hat{y}^5_{iw_i} \leq 0 && \forall i \in M \\
& -\hat{y}^1_{i0}\,(p_{i1} - p_{i0}) + \hat{y}^6_{i0} + \hat{y}^7_{i0} \leq 0 && \forall i \in M \\
& -\hat{y}^1_{ij}\left(p_{i(j+1)} - p_{ij}\right) - \hat{y}^2_{i(j-1)}\left(p_{ij} - p_{i(j-1)}\right) + \hat{y}^6_{ij} + \hat{y}^7_{ij} \leq 0 && \forall i \in M, \forall j = 1, \ldots, w_i - 1 \\
& \sum_{j=0}^{w_i} \hat{y}^3_{ij} \leq 1 && \forall i \in M \\
& \sum_{j=0}^{w_i} p_{ij}\hat{y}^3_{ij} \leq \sum_{t \in \mathcal{T}_i} \sum_{l \in \mathtt{f}(t)} c_{lt} v_t q_{lt} && \forall i \in M \\
& \sum_{k \in N} \sum_{v \in \bar{L}_k} a_{k(v-1)} \bar{x}_{kv} \leq b && \\
& \sum_{v \in \bar{L}_k} \bar{x}_{kv} = 1 && \forall k \in N \\
& \sum_{l \in \mathtt{f}(t): J^k_{lt} \subseteq \{j_1, \ldots, j_2\}} q_{lt} \leq \sum_{v=j_1}^{j_2} \bar{x}_{kv} && \begin{array}{l} \forall t \in \mathcal{T}_i, \forall i \in M, \\ \forall k \in N, \forall j_1 \leq j_2 \in \bar{L}_k \end{array} \\
& \sum_{l \in \mathtt{f}(t)} q_{lt} = 1 && \forall t \in \mathcal{T}_i, \forall i \in M
\end{aligned}
$$

$$\hat{\mathbf{y}}_i^1, \hat{\mathbf{y}}_i^5, \hat{\mathbf{y}}_i^7 \leq 0 \qquad \forall i \in M$$

$$\hat{\mathbf{y}}_i^2, \hat{\mathbf{y}}_i^3, \hat{\mathbf{y}}_i^4, \hat{\mathbf{y}}_i^6 \geq 0 \qquad \forall i \in M$$

$$\mathbf{q}_{\cdot t} \in \{0,1\}^{|\mathtt{f}(t)|} \qquad \forall t \in \mathcal{T}_i, \forall i \in M$$

$$\bar{\mathbf{x}}_k \in \{0,1\}^{|\bar{L}_k|} \qquad \forall k \in N.$$

OS.1D Tuning and Evaluating the Prediction Accuracy of *Tree Ensembles*

In this section, we discuss the tuning of TEs specifically for the collected data. Subsequently, we evaluate how well these models predict ORDs and CIs in WV.

Table OS.1 presents the tuned hyper-parameters for predicting ORDs and CIs for years 2020, 2022, and 2023. We tune four key TE parameters: (i) the maximum depth of each tree (*max_depth*), (ii) the minimum number of samples required at a leaf node (*min_samples_leaf*), (iii) the minimum number of samples needed to split an internal node (*min_samples_split*), and (iv) the number of trees in the ensemble (*n_estimators*).

Year	*max_depth*		*min_samples_leaf*		*min_samples_split*		*n_estimators*	
	ORDs	CIs	ORDs	CIs	ORDs	CIs	ORDs	CIs
2020	None	20	1	2	2	7	200	50
2022	None	30	2	2	2	5	200	100
2023	10	None	2	4	2	2	100	200

Table OS.1: Tuned TE hyper-parameters for predicting ORDs and CIs for years 2020, 2022, and 2023.

Next, we assess the predictive accuracy of TEs by comparing their performance to three commonly used regression models: *k-Nearest Neighbors* (kNN), *Support Vector Regression* (SVR), and *Linear Regression* (LR). Benchmarking TEs against LR shows their advantage over linear models, while comparisons with kNN and SVR highlight differences from instance-based and kernel-based methods in capturing non-linear patterns (İskenderoğlu et al. 2020). We assess these models over three years: 2020, 2022, and 2023. Specifically, we train the models using data up to 2019, 2020, and 2022, then predict the values of ORDs and CIs for all 55 counties in WV for the years 2020, 2022, and 2023, respectively. Year 2021 is excluded from our analysis due to data inconsistencies caused by COVID-19.

Table OS.2 compares the four regression models in predicting ORDs and CIs using three performance metrics: Coefficient of Determination (R^2), Mean Absolute Error (MAE), and a custom metric which we define as the Prediction Deviation Index (PDI). Coefficient of Determination, $R^2 \in [0,1]$, indicates how well a model explains the variability in the outcome variable (James et al. 2013); values

closer to 1 suggest that the model accounts for a large portion of the variance. MAE quantifies the average absolute difference between predicted and actual values, providing a direct measure of prediction error. The Prediction Deviation Index for model r, output i and year j is calculated as the geometric mean of $\texttt{PDI}^r_{lji} = \left|1 - \frac{P^r_{lji}}{A_{lji}}\right|$ across all 55 counties, where P^r_{lji} is the predicted value for output i in county l and year j, and A_{lji} is the corresponding observed value. While MAE captures the absolute magnitude of errors, PDI reflects the relative deviation of predictions from actual values, with smaller values indicating better accuracy. Note that $\texttt{PDI}^r_{lji}$ is undefined when $A_{lji} = 0$. In our dataset covering 2020, 2022, and 2023, there are no zero-valued CIs, and only 7 out of 165 observations for ORDs (55 counties $\times$ 3 years) have $A_{lji} = 0$. These few instances are excluded from the analysis in this section.

Year	Model	R^2		MAE		PDI	
		ORDs	CIs	ORDs	CIs	ORDs	CIs
2020	TEs	0.810	0.947	7.772	69.809	0.3241	0.2369
	kNN	0.778	0.798	8.451	99.291	0.4130	0.3096
	LR	0.775	0.935	8.854	93.226	0.4644	0.4155
	SVR	0.701	0.923	9.773	78.989	0.4880	0.3137
2022	TEs	0.919	0.906	6.170	76.779	0.3022	0.2694
	kNN	0.895	0.921	7.034	72.411	0.3882	0.3445
	LR	0.954	0.910	5.255	85.116	0.3242	0.4377
	SVR	0.940	0.865	5.413	73.215	0.3187	0.2471
2023	TEs	0.952	0.968	4.049	47.324	0.2019	0.1321
	kNN	0.923	0.904	5.942	79.376	0.3256	0.2628
	LR	0.948	0.969	4.177	68.301	0.1946	0.3794
	SVR	0.950	0.966	4.157	53.887	0.1835	0.2017

Table OS.2: R^2, MAE, and PDI ratios for ORDs and CIs across years and models.

As shown in Table OS.2, TEs deliver the highest R^2 scores for ORDs in both 2020 and 2023, and perform strongly in 2022 with a value of 0.919. For CIs, TEs also achieve the best R^2 values in 2020 and 2023, with no score dropping below 0.906 across the three years. These results highlight the robustness of TEs in explaining variation in both ORDs and CIs. In terms of accuracy, TEs obtain the lowest MAE for both outcomes in 2020 and 2023 and perform comparably to the best model in 2022. TEs also produce the lowest or near-lowest PDI values across all three years for both ORDs and CIs. Furthermore, the average PDI for TEs consistently stays below 0.33 for each year and outcome. For example, in 2022, the mean PDI is 0.302 for ORDs and 0.269 for CIs, decreasing further to 0.202 and 0.132, respectively, in

2023. This demonstrates the high predictive accuracy of TEs.

OS.2: Supplementary Figures

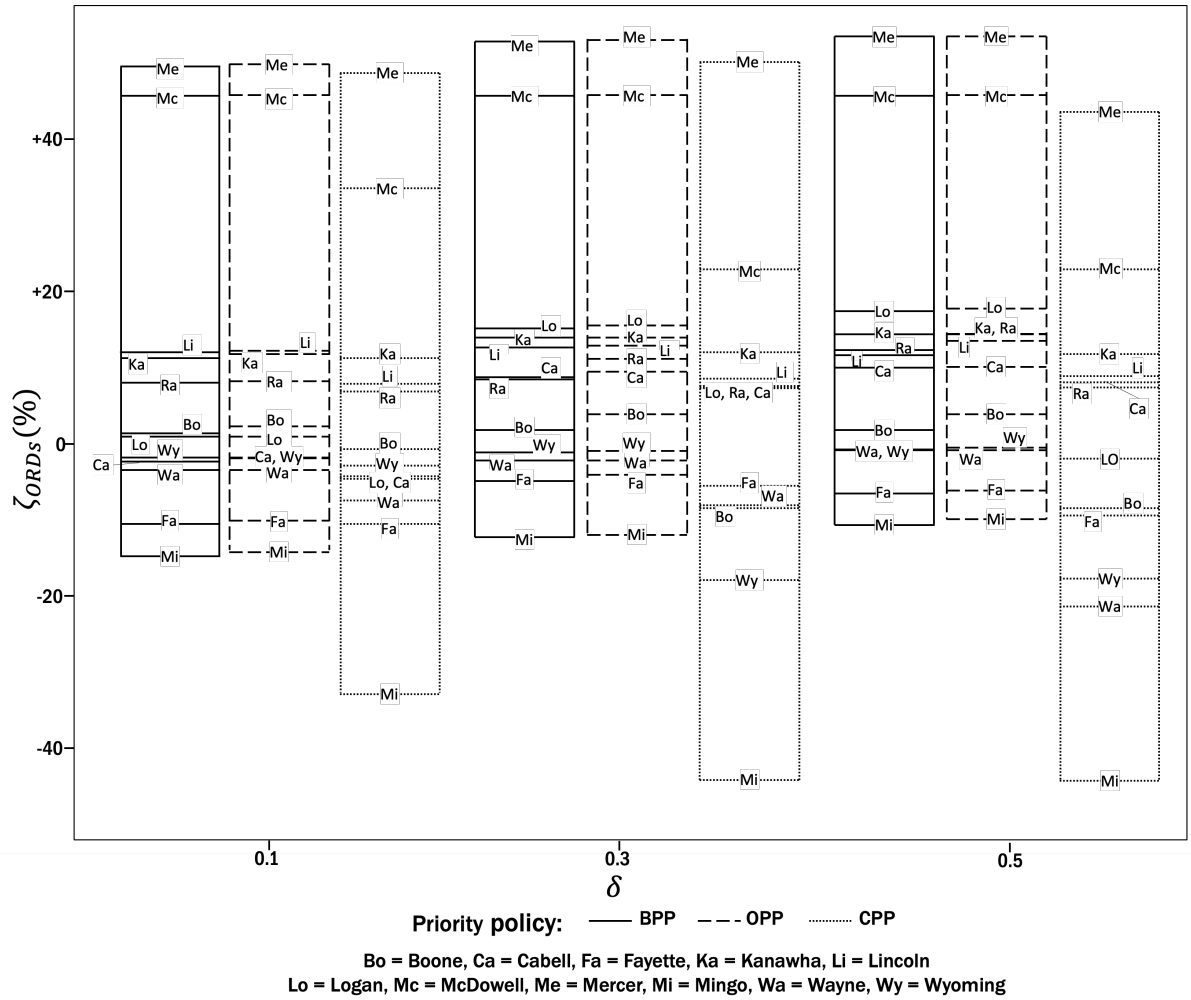

Figure OS.3: Percentage reductions in ORDs (ζ^{ORDs}) across various δ values under all *priority policies* for 12 southern counties in WV in 2023.

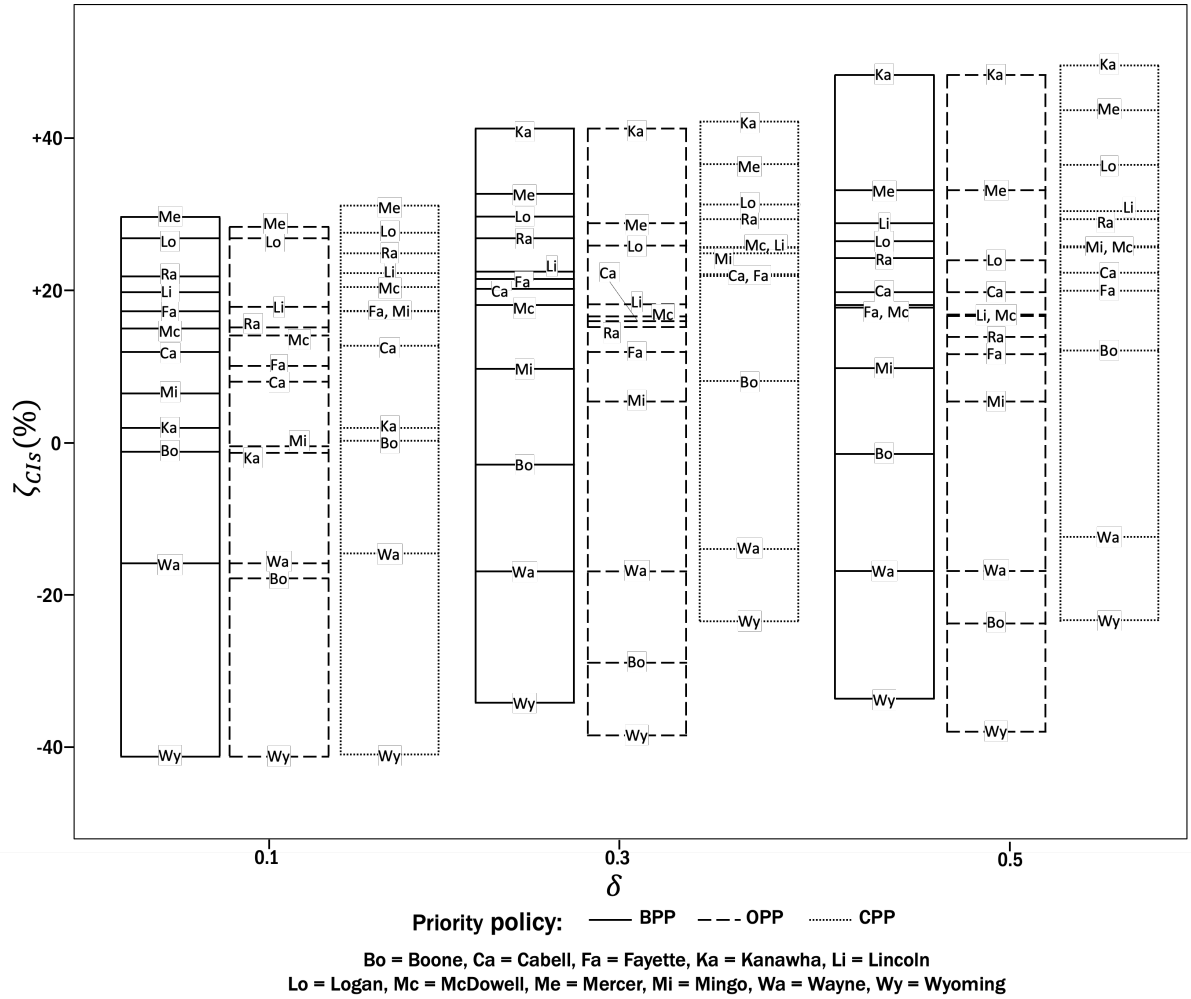

Figure OS.4: Percentage reductions in CIs ($\zeta^{\texttt{CIs}}$) across various δ values under all *priority policies* for 12 southern counties in WV in 2023.

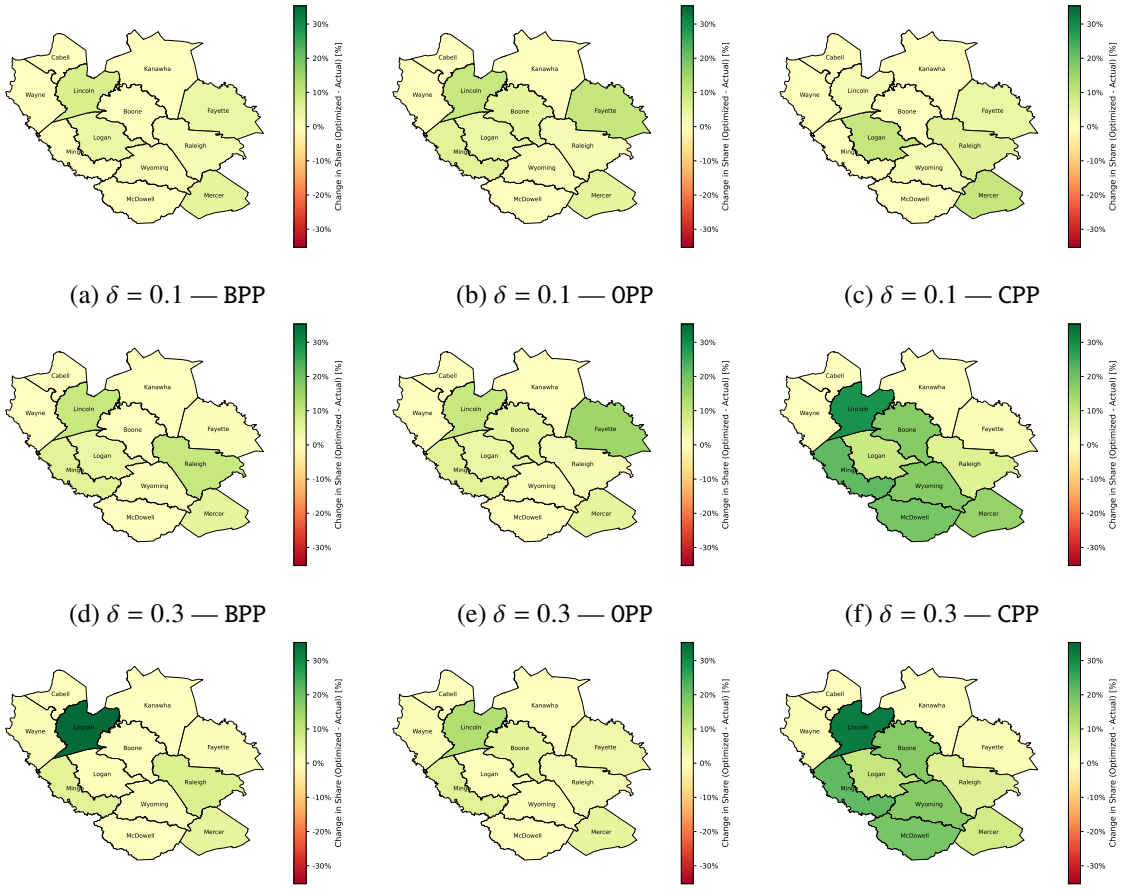

(a) $\delta = 0.1$ — BPP (b) $\delta = 0.1$ — OPP (c) $\delta = 0.1$ — CPP

(d) $\delta = 0.3$ — BPP (e) $\delta = 0.3$ — OPP (f) $\delta = 0.3$ — CPP

(g) $\delta = 0.5$ — BPP (h) $\delta = 0.5$ — OPP (i) $\delta = 0.5$ — CPP

Figure OS.5: Maximum percentage increase or minimum percentage decrease in allocation to CP ($\hat{\Delta}^{y}_{c(\mathtt{CP})}$) across various δ values under all *priority policies* for 12 southern counties in WV in 2023.

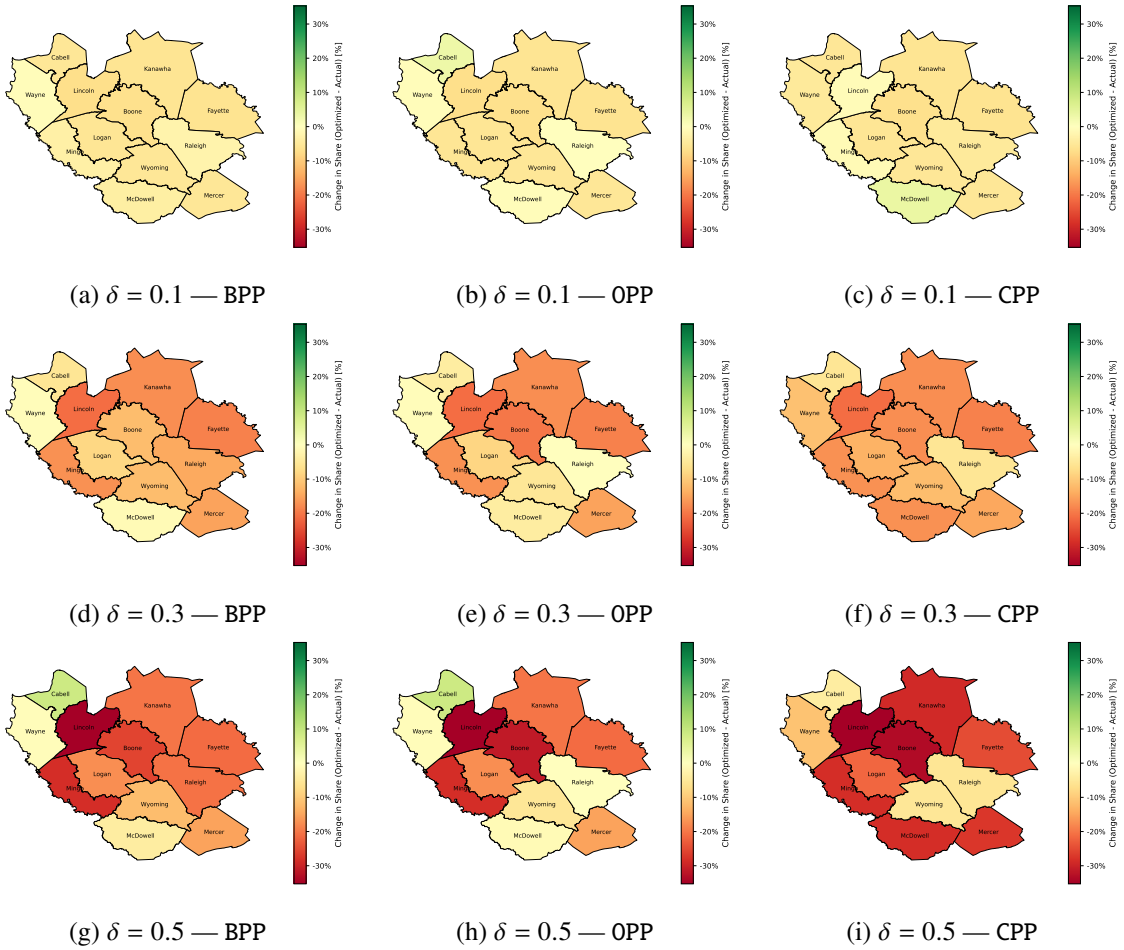

(a) $\delta = 0.1$ — BPP (b) $\delta = 0.1$ — OPP (c) $\delta = 0.1$ — CPP

(d) $\delta = 0.3$ — BPP (e) $\delta = 0.3$ — OPP (f) $\delta = 0.3$ — CPP

(g) $\delta = 0.5$ — BPP (h) $\delta = 0.5$ — OPP (i) $\delta = 0.5$ — CPP

Figure OS.6: Maximum percentage increase or minimum percentage decrease in allocation to GG ($\hat{\Delta}^{y}_{c(\texttt{GG})}$) across various δ values under all *priority policies* for 12 southern counties in WV in 2023.

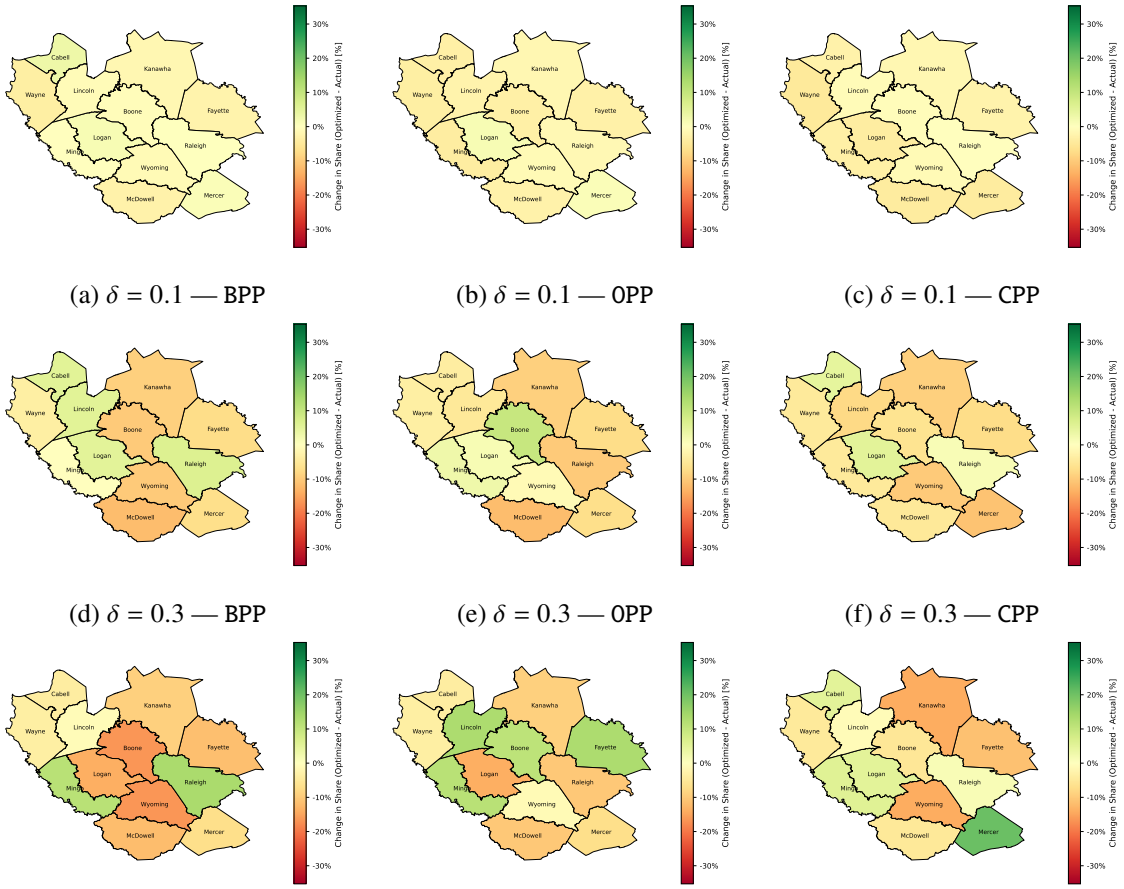

(a) $\delta = 0.1$ — BPP (b) $\delta = 0.1$ — OPP (c) $\delta = 0.1$ — CPP

(d) $\delta = 0.3$ — BPP (e) $\delta = 0.3$ — OPP (f) $\delta = 0.3$ — CPP

(g) $\delta = 0.5$ — BPP (h) $\delta = 0.5$ — OPP (i) $\delta = 0.5$ — CPP

Figure OS.7: Maximum percentage increase or minimum percentage decrease in allocation to PS ($\hat{\Delta}^{y}_{c(\text{PS})}$) across various δ values under all *priority policies* for 12 southern counties in WV in 2023.

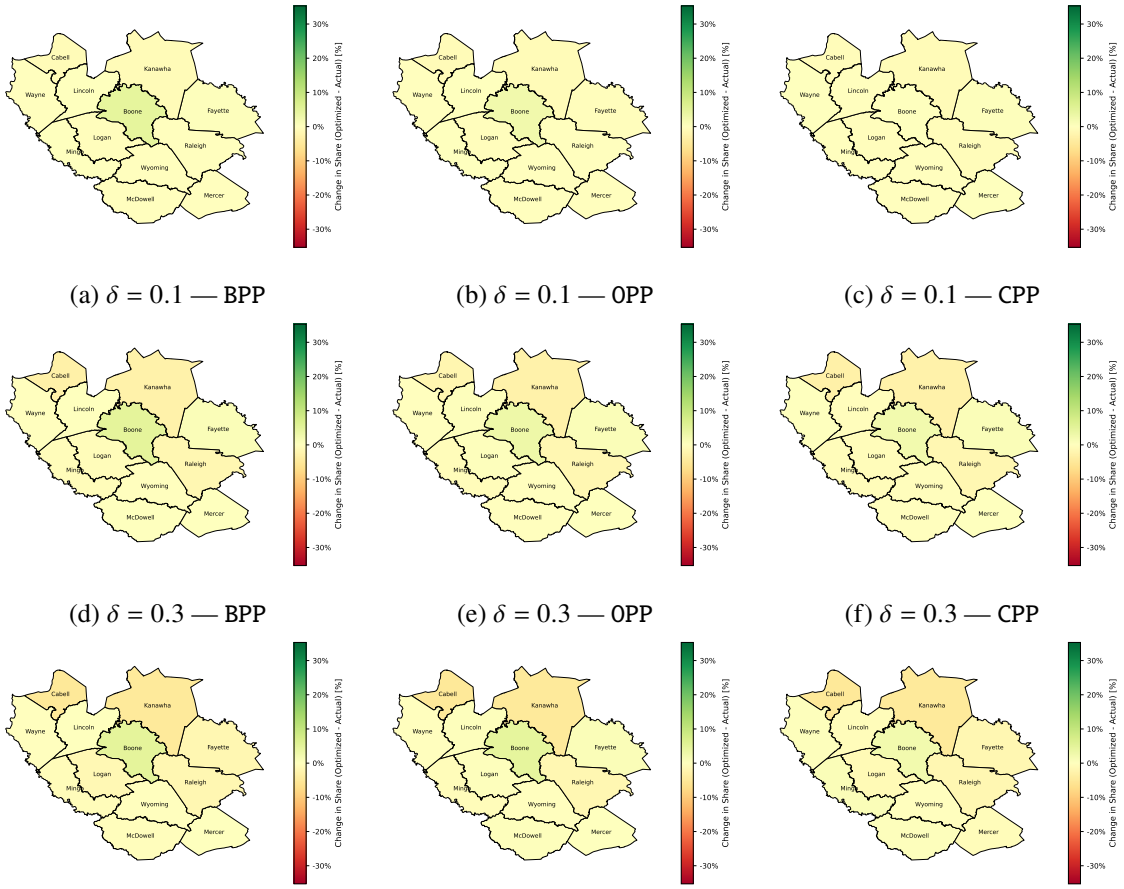

(a) $\delta = 0.1$ — BPP (b) $\delta = 0.1$ — OPP (c) $\delta = 0.1$ — CPP

(d) $\delta = 0.3$ — BPP (e) $\delta = 0.3$ — OPP (f) $\delta = 0.3$ — CPP

(g) $\delta = 0.5$ — BPP (h) $\delta = 0.5$ — OPP (i) $\delta = 0.5$ — CPP

Figure OS.8: Maximum percentage increase or minimum percentage decrease in allocation to CR ($\hat{\Delta}^{y}_{c(\mathtt{CR})}$) across various δ values under all *priority policies* for 12 southern counties in WV in 2023.

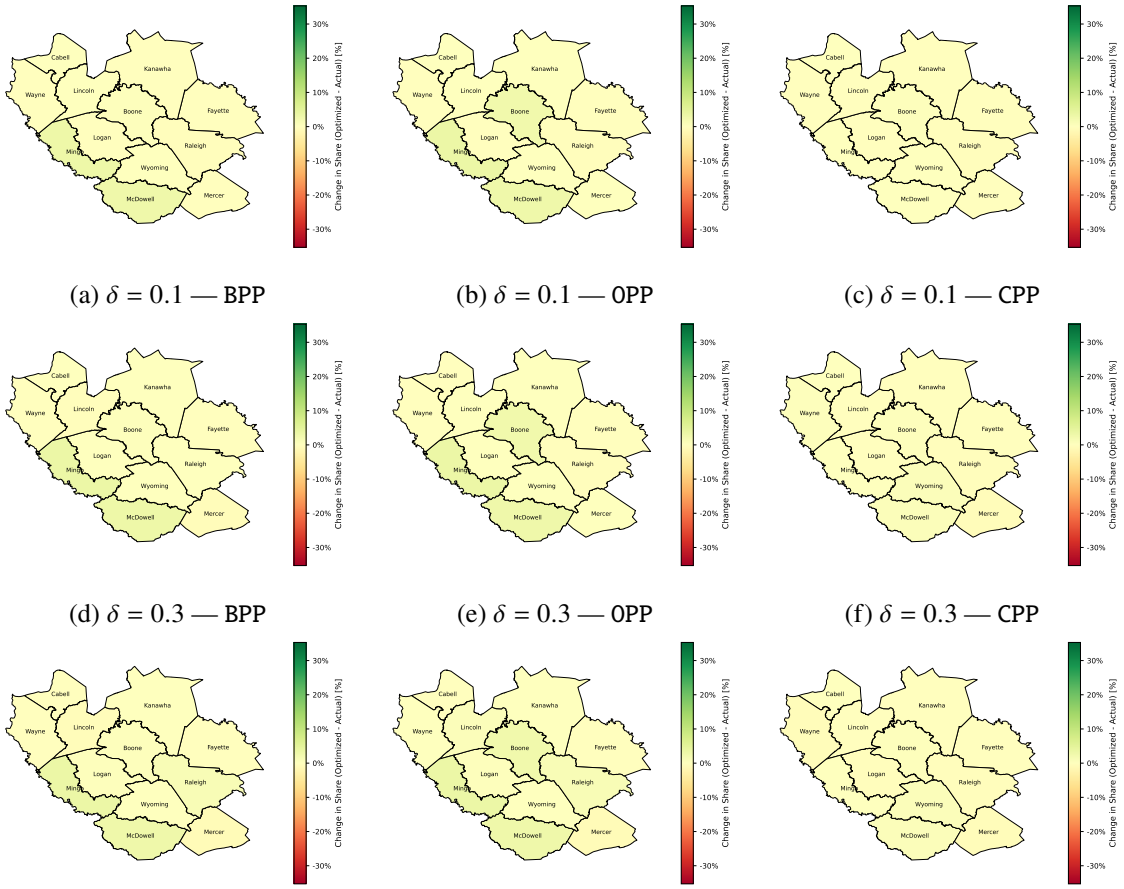

(a) $\delta = 0.1$ — BPP (b) $\delta = 0.1$ — OPP (c) $\delta = 0.1$ — CPP

(d) $\delta = 0.3$ — BPP (e) $\delta = 0.3$ — OPP (f) $\delta = 0.3$ — CPP

(g) $\delta = 0.5$ — BPP (h) $\delta = 0.5$ — OPP (i) $\delta = 0.5$ — CPP

Figure OS.9: Maximum percentage increase or minimum percentage decrease in allocation to HS ($\hat{\Delta}^{y}_{c(\texttt{HS})}$) across various δ values under all *priority policies* for 12 southern counties in WV in 2023.

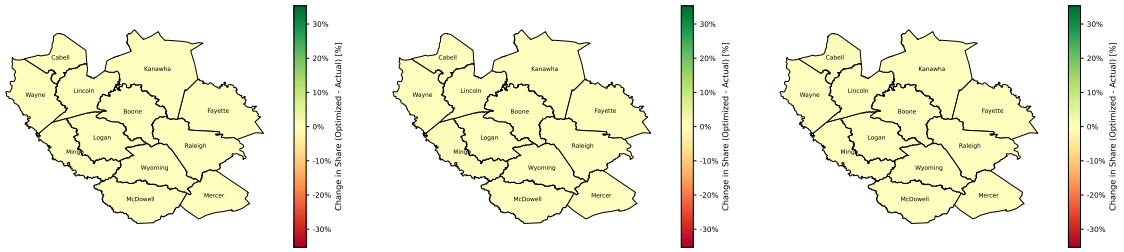

(a) $\delta = 0.1$ — BPP (b) $\delta = 0.1$ — OPP (c) $\delta = 0.1$ — CPP

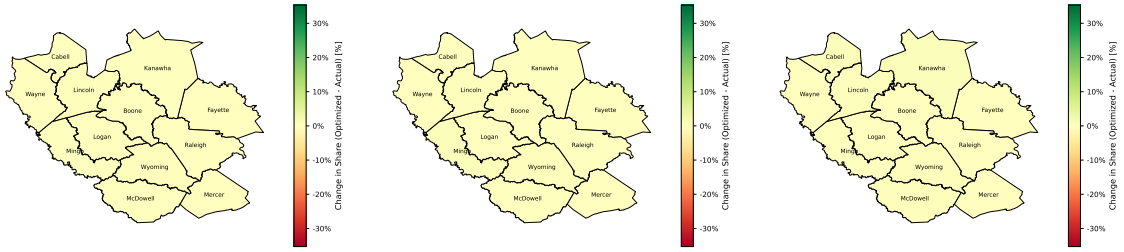

(d) $\delta = 0.3$ — BPP (e) $\delta = 0.3$ — OPP (f) $\delta = 0.3$ — CPP

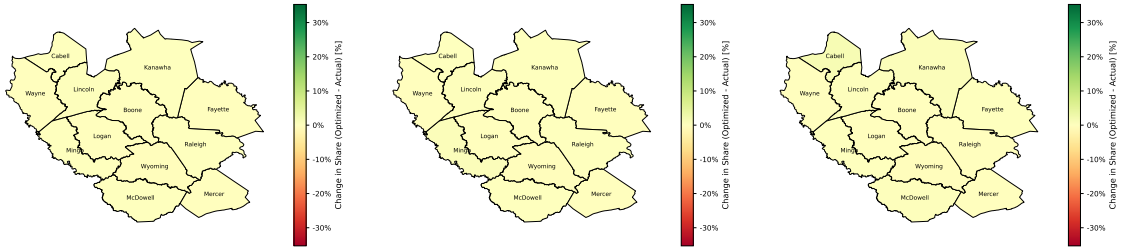

(g) $\delta = 0.5$ — BPP (h) $\delta = 0.5$ — OPP (i) $\delta = 0.5$ — CPP

Figure OS.10: Maximum percentage increase or minimum percentage decrease in allocation to SS ($\hat{\Delta}^{y}_{c(\mathtt{SS})}$) across various δ values under all *priority policies* for 12 southern counties in WV in 2023.

OS.3: Supplementary Tables

Socio-economic Variables	**Source**
Population	ACS
Uninsured (%) — Health Care Access	ACS
Median household income ($) — Income	ACS
Median gross rent ($) — Housing	ACS
Renters paying >30% of income for housing (%) — Housing	ACS
Homeowners paying >30% of income for housing (%) — Housing	ACS
Unemployment rate (% of those aged 16+) — Employment	ACS
Households receiving public assistance (%) — Food Security	ACS

Table OS.3: Socio-economic variables and their data sources.

County	CP	CR	GG	HS	PS	SS
Boone	0.00	0.00	66.89	0.00	33.11	0.01
Cabell	0.21	10.77	55.02	0.20	33.81	0.00
Fayette	0.00	5.26	63.37	0.06	31.12	0.18
Kanawha	0.17	9.80	58.81	0.62	30.59	0.00
Lincoln	0.00	1.54	70.61	1.22	26.63	0.00
Logan	0.00	3.12	56.34	0.93	39.30	0.32
McDowell	0.00	0.68	57.35	0.00	41.97	0.00
Mercer	0.00	1.14	54.99	2.07	41.56	0.25
Mingo	0.00	1.90	57.03	0.00	41.08	0.00
Raleigh	0.00	7.78	48.13	3.24	40.42	0.42
Wayne	1.71	0.81	48.72	1.40	47.37	0.00
Wyoming	0.00	1.38	60.76	3.41	33.35	1.10

Table OS.4: Expenditure share (%) by category for each county in 2023.