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`), a state severely affected by the epidemic. Drawing on this data set, we employ a Fixed Effects (`FE`) model to estimate the relationships between county-level budget allocations and two outcomes associated with the opioid crisis: Opioid-Related Deaths (`ORDs`) and Crime Incidents (`CIs`). We then develop combined `FE` and Tree Ensemble (`TEs`) models, in which `TEs` capture remaining input-output relationships not explicitly represented in the linear `FE` model. We embed the resulting `FE+TEs` models within a Mixed Integer Programming (MIP) formulation that identifies budget allocations across expenditure categories to maximize policymakers' worst-case, risk-averse utility over `ORDs` and `CIs`. We apply the framework to 12 counties in southern `WV` that have been most impacted by the opioid crisis. When the two outcomes are given equal emphasis, the recommended allocations reduce the predicted levels of both `ORDs` and `CIs`; however, the results also reveal trade-offs between the two outcomes when policymakers place greater emphasis on one relative to the other. Furthermore, sensitivity analysis with respect to the total available budget shows that additional resources do not necessarily improve predicted outcomes when the flexibility to reallocate those resources across expenditure categories is limited. More broadly, the results demonstrate how the proposed framework can generate tailored, county-level budget recommendations to assist policymakers in reprioritizing expenditures to 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). In 2023

*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

alone, approximately 105,000 people died from drug overdoses, nearly 80,000 of which involved opioids, nearly ten times the number of opioid overdose deaths recorded in 1999 (Centers for Disease Control and Prevention 2025). This has raised concern among policymakers at both federal and state levels (Office of National Drug Control Policy 2022). Limiting the spread of opioid misuse and mitigating the impacts of the opioid crisis remain at the forefront of public discourse across the U.S.

While recent DEA National Threat Assessment reports (Drug Enforcement Administration 2025) provide substantial insight into addressing the opioid epidemic at the national level, effective mitigation also requires understanding how resources are allocated within local communities. Local governments play an important role in this response by allocating resources to services that directly affect their communities. Among local government entities, counties provide a meaningful unit of analysis because they administer and coordinate services central to opioid-crisis mitigation, including public health, behavioral health, emergency response, social services, and public safety (National Association of Counties and Appalachian Regional Commission 2019). Examining counties within a single state further allows us to study variation in local conditions, outcomes, and resource allocation within a common state-level policy environment.

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 2026). 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 Fixed Effects (`FE`) model (Wooldridge 2010) to estimate the impact of budget allocations across categories on outcomes associated with the opioid crisis.

We focus on West Virginia (`WV`) as a test case for our county-level decision-making framework. In 2022, `WV` had an age-adjusted drug-overdose death rate of 80.9 per 100,000 population and an opioid-involved overdose death rate of 70.5 per 100,000. These rates were substantially higher than the corresponding national rates of 32.6 and 25.0 per 100,000 population, respectively, highlighting the disproportionate impact of the opioid crisis on `WV` (Garnett and Miniño 2025). `WV` also provides an important rural context, with 55.4% of its population residing in rural areas, compared with 25.9%

nationally (U.S. Census Bureau 2023). Rural communities often face additional barriers to treatment and service access, increasing the importance of how limited local resources are allocated across services that can help address the opioid crisis. Counties play an important role in these allocation decisions through their responsibilities for health, social, emergency, and public safety services (National Association of Counties and Appalachian Regional Commission 2019). Moreover, WV exhibits substantial county-level variation in overdose outcomes and expenditures directed toward mitigating the crisis, providing a heterogeneous setting in which to develop and evaluate the proposed framework. Although this paper focuses on WV, the underlying challenge of allocating limited county resources across expenditure categories to improve opioid-related and broader community outcomes is not unique to the state. Other states with substantial opioid-related mortality, including Delaware, Tennessee, Vermont, Connecticut, and Alaska, may similarly benefit from county-level resource allocation approaches (Garnett and Miniño 2025).

While our mathematical framework is applicable to any set of outcomes important to a county, this paper focuses on two distinct dimensions of the opioid crisis: Opioid-Related Deaths (ORDs) as a direct measure of its health consequences and Crime Incidents (CIs) as a measure of potential broader community-level consequences. The inclusion of CIs is motivated by a growing empirical literature documenting links between substance misuse, opioid exposure, and criminal activity. Substance use may affect crime through several mechanisms, including pharmacological effects that influence aggressive behavior, economically motivated property crime to finance drug consumption, and violence associated with participation in illicit drug markets (Bondurant et al. 2018, Dave et al. 2021). Consistent with these mechanisms, Bondurant et al. (2018) find that increased access to substance-abuse treatment facilities reduces local violent and financially motivated crimes, while Dave et al. (2021) find that mandatory-access prescription drug monitoring programs reduce overall crime, with reductions particularly in aggravated assault and burglary. More recently, Sim (2023) finds that greater geographic exposure to OxyContin was associated with higher violent crime rates, further supporting the consideration of crime as a broader community-level outcome associated with opioid exposure. Taken together, these studies suggest that the consequences of opioid misuse and efforts to address it may extend beyond direct health outcomes to criminal activity. At the same time, these findings do not imply that criminal activity is attributable solely to opioid misuse, as CIs represent a broader community-level outcome that may arise through both opioid-related and other pathways.

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 an analytical tool that is meant to supplement, and not replace, exist-

ing 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. Developing such a framework requires addressing three key considerations:

(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 address (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 (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 address (c), we develop a combined FE and *Tree Ensembles* (TEs) framework to characterize the relationships between budget allocations and both ORDs and CIs. We characterize these relationships without explicitly modeling the specific pathways through which expenditures may affect ORDs and CIs, as comprehensive county-level data on these intermediate mechanisms is unavailable. The FE component estimates the effects of explanatory variables on outcomes while accounting for unobserved time-invariant county characteristics and common year-specific effects. We then employ TEs to capture additional input–output relationships that are not explicitly represented in the linear FE model. Following a predict-then-optimize approach, we integrate the trained FE+TEs models into a Mixed Integer Programming (MIP) formulation to identify budget allocations that improve the considered outcomes subject to the county's available budget and other allocation constraints. Subsequently, we analyze the resulting budget allocation recommendations for selected counties in WV that have been impacted

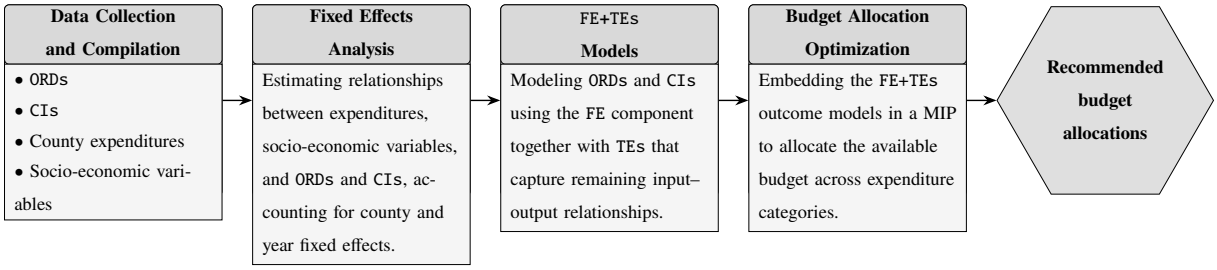


Figure 1: Overview of the proposed data-driven decision-support framework for county-level budget allocation.

the most by the opioid crisis. Figure 1 provides an overview of the proposed data-driven decision-support framework and illustrates how county-level data, FE model, FE+TEs models, and optimization are integrated to generate budget allocation recommendations.

Our empirical and computational analyses yield several key findings. The FE models identify significant relationships between county expenditures and both ORDs and CIs. Using the combined FE+TEs models, the proposed optimization framework identifies county-specific budget allocations that can improve predicted ORDs and CIs. The results also reveal important trade-offs between the two outcomes, as allocations that generate greater improvements in one outcome may produce smaller improvements or deterioration in the other. Finally, our sensitivity analysis shows that increasing the available budget does not necessarily improve predicted outcomes when constraints on deviations from existing budget patterns limit how additional resources can be allocated across 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 presents the FE models of the relationships between budget allocations, socio-economic variables, and ORDs and CIs. Section 5 introduces the proposed budget allocation model. Section 6 analyzes the model results. Lastly, Section 7 offers concluding remarks. This paper is accompanied by an Online Supplement, where the labels for Sections, Tables, and Figures are prefixed with OS.

2 Literature Review

Computational research on the opioid epidemic has frequently used compartmental epidemic models to characterize its progression and dynamics (White and Comiskey 2007, Battista et al. 2019). These models define a set of states associated with different stages of opioid use and allow individuals to transition between these states over time. 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 Susceptible-Infected-Removed (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 the `FE+TE`-based predictive models to estimate `ORDs` and `CIs` as functions of county-level budget allocations across various expenditure categories. Table 1 summarizes the computational modeling and optimization studies related to the opioid epidemic.

More broadly, mathematical optimization has been widely used to support healthcare resource allocation under limited budgets and capacity constraints (Rais and Viana 2011). These models provide a systematic way of allocating scarce resources while accounting for competing objectives, intervention interactions, and operational constraints. More recent studies have combined data-driven prediction with optimization to support healthcare resource planning. For example, Mizan and Taghipour (2022) integrate machine-learning forecasts of patient demand with multi-objective optimization for radiological resource allocation, while Olya et al. (2022) combine deep learning with stochastic optimization for healthcare workforce planning. In the context of the opioid crisis, data-driven optimization has also been used to allocate specific interventions across geographic regions, including treatment facilities (Luo and Stellato 2024, Baucum et al. 2025) and medications for opioid use disorder, such as buprenorphine (Dong et al. 2023). Our work differs from these studies by considering the allocation of a county's budget across multiple expenditure categories rather than allocating a single healthcare resource or intervention.

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

Study	**Estimation approach**	**Decision/Intervention**	**Outcome/Objective**	**Optimization approach**
White and Comiskey (2007)	ODE-based epidemic model	Treatment and prevention interventions	Heroin-use dynamics	–
Battista et al. (2019)	Compartmental epidemic model	Prescription-opioid dynamics	Opioid-use progression	–
Chen et al. (2019)	Three-compartment system dynamics model	Prescription-opioid misuse interventions	Opioid-related deaths	–
Pitt et al. (2018)	Dynamic compartmental model	Alternative opioid-policy interventions	Deaths, life years, and quality-adjusted life years	–
Rao et al. (2021)	Dynamic simulation model	Alternative opioid-policy interventions	Opioid-related health outcomes	–
Ansari et al. (2024)	SIR compartmental model	Budget allocation across preventive and mitigating interventions	Economic costs associated with fatal overdoses	Markov Decision Process
Luo and Stellato (2024)	Differential-equation-based epidemic model	Treatment-facility locations and treatment budgets	Opioid-related deaths and population with opioid use disorder	Mixed-integer optimization
Baucum et al. (2025)	Two-stage least squares	Allocation of SUD treatment centers across counties	Overdose deaths and geographic disparities	Mixed-integer quadratic programming
Dong et al. (2023)	Bayesian spatio-temporal model	County-level buprenorphine allocation	Illicit-opioid overdose mortality	Constrained nonlinear optimization
This study	`FE+TEs`	County-level budget allocation	`ORDs` and `CIs`	MIP with robust risk-averse utility

Table 1: Summary of computational modeling and optimization studies related to the opioid epidemic.

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. 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`.

Our work builds on this predict-then-optimize literature but differs in both the estimation and decision-making components. We develop a combined `FE+TEs` model in which the `FE` component estimates the relationships between county budget allocations, socio-economic variables, and the outcomes while accounting for county and year fixed effects. The `TE` component captures input–output relationships not represented by the linear `FE` model. We embed the resulting outcome models within an MIP formulation that reallocates a county budget across multiple expenditure categories. The formulation simultaneously considers two outcomes, `ORDs` and `CIs`, through a robust risk-averse utility maximization framework. Thus, our main contribution to the data-driven resource allocation literature is a framework that integrates `FE+TEs` models with a robust risk-averse optimization approach that accounts for multiple outcomes and unknown decision-maker preferences when allocating county budgets across multiple expenditure categories. We apply the proposed framework to 12 counties in `WV` that have been impacted the most by the opioid crisis and analyze the resulting budget allocation recommendations. Although the application focuses on `ORDs` and `CIs`, the framework can accommodate alternative outcomes and decision-maker preferences, providing a general approach for data-driven county-level resource allocation.

3 Data Collection and Compilation

We collect data from all 55 counties in `WV` for fiscal years (FY) 2012-2023. We use data from 2012–2022 to develop our proposed budget allocation model and reserve 2023 as an out-of-sample year for model evaluation and application of the proposed budget allocation framework.

We focus on two opioid epidemic-related outcomes: `ORDs` and `CIs` in `WV`. For `ORDs`, we include all drug-overdose fatalities in which opioid use was identified as a primary cause of death. 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. For CIs, we consider both property and violent crimes, as prior research identifies both as important societal outcomes associated with opioid misuse. Substance use disorders may contribute to property crime through income-generating activities used to finance drug consumption, while violent crime may arise through behavioral effects of substance misuse and interactions within illicit drug markets (Dave et al. 2021, Sim 2023). CIs are measured using data from the West Virginia Incident-Based Reporting System (WVIBRS), a state-level implementation of the FBI's National Incident-Based Reporting System (NIBRS), and consist of county-level annual counts of five reported incident types: Burglary, Aggravated Assault, Simple Assault, Intimidation Assault, and Robbery. Figure OS.1 presents the distributions of the observed and log-transformed ORDs and CIs per 100,000 population across all 55 WV counties over the 2012–2022 period. To accommodate zero values, we add 1 to each normalized observation before applying the log-transformation. As shown in the figure, the original distributions of both outcomes exhibit substantial right-skewness, with a small number of relatively large observations. We therefore use log-transformed outcomes in our analysis to reduce skewness and mitigate the influence of extreme values.

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. 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 2026, 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*)*. Based on their primary functions and the types of services they support, we aggregate the six expenditure categories into three groups: CP + HS, CR + GG + PS, and SS. This grouping reflects three distinct but complementary functions of county government: CP + HS primarily captures expenditures aimed at developing and maintaining a county's physical infrastructure and public facilities, while supporting community health through healthcare services, and related infrastructure; CR + GG + PS represents ex-

penditures associated with maintaining government operations, public order, and community well-being through administrative services, law enforcement, corrections, public safety, and community programs; and SS primarily captures expenditures directed toward meeting residents' socioeconomic needs and providing a safety net through services such as human resources, public transit, homeless shelters, and family support programs. We retain SS as a separate category because these expenditures focus more directly on individual and household support, distinguishing them from the infrastructure- and health-related functions of CP + HS and the governmental, community, and public-safety functions of CR + GG + PS. Throughout the remainder of the paper, we refer to these three aggregated expenditure categories as AECs. Figure OS.2 presents the distributions of AECs per 100,000 population in our dataset before and after the log-transformation. Similar to ORDs and CIs, we add 1 to each normalized AEC observation before applying the log-transformation to accommodate zero values. As shown in the figure, the original expenditure distributions exhibit substantial right-skewness, with a small number of relatively large observations. Therefore, we use log-transformed expenditures in our analysis to reduce skewness and mitigate the influence of extreme values.

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. 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 four socio-economic variables, with one representative variable for each of these groups. The uninsured rate (uninsured_pct) represents Health Care Access and measures the percentage of people without health insurance coverage. Housing cost burden (own30_pct) represents Housing and measures the percentage of owner-occupied housing units with housing costs exceeding 30% of household income. Median household income (med_hh_inc) represents Employment and measures the median annual household income in the county, in dollars. Public assistance receipt (pub_assist_pct) is used as a proxy for Food Security (Cesare et al. 2024) and measures the percentage of households receiving public cash assistance during the past 12 months. We also incorporate each county's population as an additional input in our model. Population (pop) is the number of residents in the county, measured in persons. For ease of exposition, we categorize population alongside the socio-economic variables. Table OS.1 presents the data sources and descriptive statistics (mean, minimum, maximum, and standard deviation) for the % of total budget share for each AECs and considered socio-economic variables across 605 county-year observations from 55 West Virginia counties during 2012–2022. Furthermore, the interested reader will find maps illustrating county-level data for budget, ORDs, CIs, and population in Figure OS.3.

4 Fixed Effects Models

In this section, we employ the Fixed Effects (FE) model (Wooldridge 2010) to estimate the 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 11-year span from 2012 to 2022. The FE models for ORDs and CIs in county c during year t are specified as follows:

$$\overline{\texttt{ORDs}}_{ct} = \sum_{j \in J} \beta_j^{\texttt{ORDs}} x_{cj(t-1)} + \sum_{s \in S_e} \gamma_s^{\texttt{ORDs}} h_{cs(t-1)} + \alpha_c^{\texttt{ORDs}} + \kappa_t^{\texttt{ORDs}} + \varepsilon_{ct}^{\texttt{ORDs}} \tag{1a}$$

$$\overline{\texttt{CIs}}_{ct} = \sum_{j \in J} \beta_j^{\texttt{CIs}} x_{cj(t-1)} + \sum_{s \in S_e} \gamma_s^{\texttt{CIs}} h_{cs(t-1)} + \alpha_c^{\texttt{CIs}} + \kappa_t^{\texttt{CIs}} + \varepsilon_{ct}^{\texttt{CIs}}. \tag{1b}$$

Here, $\overline{\texttt{ORDs}}_{ct}$ and $\overline{\texttt{CIs}}_{ct}$ denote the log-transformed ORDs and CIs per 100,000 population, respectively, in county c and year t. J denotes the set of expenditure categories, and S_e represents the set of socio-economic variables. The variable $x_{cjt} = \log(1 + \tilde{x}_{cjt})$ where $\tilde{x}_{cjt}$ indicates the budget amount allocated per 100,000 people by county c to expenditure category j in year t. The variable h_{cst} represents the value of socio-economic variable s for county c in year t. We use one-year-lagged explanatory variables so that budget allocations and socio-economic conditions are measured one year before the corresponding outcomes, allowing for potential delays in their effects on opioid-related outcomes. The coefficients $\beta_j^{\texttt{ORDs}}$ and $\beta_j^{\texttt{CIs}}$ capture the associations between lagged budget allocations and ORDs and CIs, respectively, while $\gamma_s^{\texttt{ORDs}}$ and $\gamma_s^{\texttt{CIs}}$ capture the corresponding effects of lagged socio-economic variables. The terms $\alpha_c^{\texttt{ORDs}}$ and $\alpha_c^{\texttt{CIs}}$ denote county-specific fixed effects, while $\kappa_t^{\texttt{ORDs}}$ and $\kappa_t^{\texttt{CIs}}$ denote year fixed effects for the ORDs and CIs models, respectively. Finally, $\varepsilon_{ct}^{\texttt{ORDs}}$ and $\varepsilon_{ct}^{\texttt{CIs}}$ are the idiosyncratic error terms for the respective models.

The causal interpretation of the FE models in (1) relies on the strict exogeneity assumption, conditional on the county and year fixed effects as well as the included time-varying explanatory variables (Assumption FE.1, Page 301, Wooldridge (2010)). In simple terms, this assumption requires that, after accounting for the included explanatory variables and fixed effects, unobserved factors affecting the outcome in a given year are not associated with past, current, or future values of the explanatory variables. Several features of our FE model support the plausibility of this assumption. County fixed effects account for time-invariant observed and unobserved heterogeneity across counties that may be correlated with the considered outcomes, whereas year fixed effects absorb common temporal shocks affecting all counties. In addition, budget allocations and socio-economic variables enter the FE models (1) with a one-year lag. This specification establishes temporal ordering between the explanatory variables and the outcomes, thereby reducing concerns about contemporaneous reverse causality. To further assess the strict exogeneity

assumption, we conduct a lead-variable test (Wooldridge 2010, Su et al. 2016). Specifically, we augment the FE models in (1) with one-period-ahead values of the explanatory variables and conduct a joint F-test of the null hypothesis that the coefficients of the future explanatory variables are jointly equal to zero. For ORDs, the test yields an F-statistic of 0.0741 with a p-value of 0.9737, while for CIs, the corresponding F-statistic is 0.6211 with a p-value of 0.6044. Because both p-values significantly exceed the 0.10 significance level, we fail to reject the null hypothesis for either outcome. This indicates that the future budget allocation variables do not provide statistically significant additional explanatory power after accounting for the included explanatory variables, and county and year fixed effects.

Explanatory Variable	VIF
$\log(\text{CP} + \text{HS})_{t-1}$	1.1024
$\log(\text{CR} + \text{GG} + \text{PS})_{t-1}$	1.2568
$\log(\text{SS})_{t-1}$	1.1774
$\texttt{own30_pct}_{t-1}$	1.2689
$\texttt{med_hh_inc}_{t-1}$	1.3496
$\texttt{pub_assist_pct}_{t-1}$	1.0625
$\texttt{uninsured_pct}_{t-1}$	1.0858
$\texttt{pop}_{t-1}$	1.1431

Table 2: Variance Inflation Factors (VIF) for the input variables

We also assess potential multicollinearity among the explanatory variables included in the FE models using variance inflation factors (VIFs) (Hair et al. 2022). Table 2 reports the resulting VIFs. All VIF values are close to one, ranging from 1.0625 to 1.3496. These low VIF values indicate limited dependence among the explanatory variables after accounting for county and year fixed effects, suggesting that multicollinearity is unlikely to affect the estimated coefficients in the FE models.

Tables 3 and 4 present the results of the FE models in (1a) and (1b), respectively, obtained using *Stata 19* (StataCorp 2025) and the panel dataset described in Section 3. The first column lists the explanatory variables, while the subsequent columns report the estimated coefficients, their corresponding p-values, and the associated 95% confidence intervals. These results reveal consistent patterns across ORDs and CIs for the AECs. The coefficient of $\log(\text{CP}+\text{HS})_{t-1}$ is negative for both outcomes, with statistical significance at approximately 5% level. These results suggest that higher per-capita expenditures on CP and HS, which encompass investments in physical infrastructure, public facilities, healthcare, and related services, are associated with lower subsequent ORDs and CIs within a county. In contrast, $\log(\text{CR} + \text{GG} + \text{PS})_{t-1}$ has a positive and statistically significant coefficient for both ORDs and CIs, indicating a positive association

Explanatory Variable	Coeff.	p-value	95% Confidence Interval
$\log(\text{CP} + \text{HS})_{t-1}$	-0.018^{*}	0.055	$[-0.037, 0.0004]$
$\log(\text{CR} + \text{GG} + \text{PS})_{t-1}$	0.420^{*}	0.088	$[-0.065, 0.905]$
$\log(\text{SS})_{t-1}$	0.012	0.450	$[-0.019, 0.042]$
$\texttt{own30_pct}_{t-1}$	0.029	0.349	$[-0.032, 0.090]$
$\texttt{med_hh_inc}_{t-1}$	-1.5×10^{-5}	0.321	$[-4.5 \times 10^{-5}, 1.5 \times 10^{-5}]$
$\texttt{pub_assist_pct}_{t-1}$	0.116^{**}	0.036	$[0.008, 0.224]$
$\texttt{uninsured_pct}_{t-1}$	0.021	0.657	$[-0.075, 0.118]$
$\texttt{pop}_{t-1}$	2.1×10^{-5}	0.290	$[-1.8 \times 10^{-5}, 6.0 \times 10^{-5}]$

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

Table 3: Results of the `FE` model for `ORDs`.

Explanatory Variable	Coeff.	p-value	95% Confidence Interval
$\log(\text{CP} + \text{HS})_{t-1}$	-0.024^{**}	0.048	$[-0.048, -0.002]$
$\log(\text{CR} + \text{GG} + \text{PS})_{t-1}$	0.424^{***}	0.001	$[0.175, 0.674]$
$\log(\text{SS})_{t-1}$	-0.008	0.212	$[-0.022, 0.005]$
$\texttt{own30_pct}_{t-1}$	-0.0111	0.4226	$[-0.0385, 0.0164]$
$\texttt{med_hh_inc}_{t-1}$	-9.9×10^{-6}	0.3154	$[-2.9 \times 10^{-5}, 9.6 \times 10^{-6}]$
$\texttt{pub_assist_pct}_{t-1}$	-0.0277	0.4453	$[-0.100, 0.045]$
$\texttt{uninsured_pct}_{t-1}$	-0.049^{*}	0.054	$[-0.098, 0.001]$
$\texttt{pop}_{t-1}$	$-3.5 \times 10^{-5***}$	0.0018	$[-5.6 \times 10^{-5}, -1.3 \times 10^{-5}]$

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

Table 4: Results of the `FE` model for `CIs`

between expenditures in this category and both subsequent outcomes. Given that CR+GG+PS accounts for the vast majority of county expenditures, these findings highlight the importance of allocation trade-offs under finite county budgets. Higher expenditures on CR + GG + PS may coincide with lower expenditures on other categories, including CP and HS, which exhibit negative associations with both outcomes. These contrasting relationships suggest that reallocating expenditures from CR + GG + PS toward CP + HS may offer opportunities to reduce predicted ORDs and CIs without increasing the total county budget.

The results for the socio-economic variables differ more noticeably between the two outcomes. For ORDs, the percentage of individuals receiving public assistance has a positive and statistically significant coefficient. This positive association reflects broader underlying socioeconomic vulnerabilities that are associated with an elevated risk of opioid-related mortality (Cesare et al. 2024). For CIs, the percentage of the population without health insurance has a negative and statistically significant coefficient. One possible interpretation for this association is that the uninsured rate may capture dimensions of socioeconomic disadvantage that are not fully represented by the other explanatory variables in the model. Previous research suggests that socioeconomic inequalities and marginalization can contribute to distrust and weaker engagement between communities and law enforcement (Reinsch and Xiu 2025). Because CIs are derived from NIBRS data, they capture incidents reported to and recorded by participating law-enforcement agencies rather than all criminal activity occurring within a county. Thus, the negative coefficient may reflect lower reporting and recording of crime in counties with higher uninsured rates. A similar mechanism may also explain the statistically significant negative coefficient for population, as more populous counties in WV tend to have a larger number of socioeconomically disadvantaged residents (National Institute on Minority Health and Health Disparities 2026).

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 the AECs (CP+HS,CR+GG+PS,SS) in a given year. The goal is to reduce ORDs and CIs. While our numerical experiments in Section 6 consider specifically these AECs and the two outcomes, we next derive a model for the general case involving set of expenditure categories J and set of outcomes I. This general formulation enables broader applicability across states, outcomes, and expenditure structures. We present our modeling framework for determining budget allocation of county $\tilde{c}$ in year $\tilde{t}$. For notational simplicity, we omit the indices $\tilde{c}$ and $\tilde{t}$ wherever they are unambiguous and use boldface lowercase letters to denote vectors. As in Section 4, $\mathbf{x}$ denotes the vector of log-transformed budget allocations per 100,000 population across expenditure categories for county $\tilde{c}$ in year $\tilde{t}$, while $\mathbf{h}$ denotes the vector containing the values of the socioeconomic variables for the same county and year. For a given $(\mathbf{x}, \mathbf{h})$, $\xi_i(\mathbf{x}, \mathbf{h})$ denotes the value of outcome $i \in I$. We denote the vector of

all such outcomes (learned from data) by $\bar{\xi}(\mathbf{x}, \mathbf{h}) = \big(\xi_1(\mathbf{x}, \mathbf{h}), \xi_2(\mathbf{x}, \mathbf{h}), \ldots, \xi_{|I|}(\mathbf{x}, \mathbf{h})\big)^\top$.

Using this notation, the model we solve is

$$\max_{\mathbf{x} \in \mathbf{X}} \min_{u \in \mathbf{U}} u(\bar{\xi}(\mathbf{x}, \mathbf{h})), \tag{2}$$

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

$$\mathbf{X} = \left\{ \mathbf{x} \in \mathbb{R}^{|J|} \;\middle|\; \sum_{j \in J} (e^{x_j} - 1) \le b, \qquad \log\Big(max(1, 1 + \hat{x}_j - \kappa_j^-)\Big) \le x_j \le \log(1 + \hat{x}_j + \kappa_j^+), \quad \forall j \in J \right\}. \tag{3}$$

Model (2) 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 (3) ensures that: (i) the total available budget b is not over-allocated across $|J|$ expenditure categories, and (ii) the amount allocated to each category j can decrease from the *reference budget allocation* $\hat{x}_j$ by at most κ_j^- or increase by at most κ_j^+. These latter constraints, which we refer to as *reference budget allocation* (`RBA`) constraints, restrict annual changes by keeping each expenditure category's allocation close to the previous year's value, reflecting counties' limited flexibility in adjusting budgets (Kelly 2005, Good 2011).

Although model (2) 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 develop a combined `FE+TEs` framework that integrates the relationships identified by the `FE` model in Section 4 with the predictive flexibility of `TEs`. Specifically, the resulting predictive model combines the `FE` model with a `TEs`-based correction trained on the residuals of the `FE` component to capture additional input–output relationships. In addition to their predictive accuracy, a key advantage of `TEs` in our setting is that their input-to-output relationships can be represented through mixed-integer linear constraints. This enables us to incorporate the resulting `FE+TEs` predictive function, $\bar{\xi}(\mathbf{x}, \mathbf{h})$, directly into (2). Second, in Section 5.2, we describe a set of 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 (2) can be formulated and solved as an MIP, which we describe in Section 5.3.

5.1 Modeling Outcome Functions $\xi_i(\mathbf{x}, \mathbf{h})$ using FE+TEs

We predict each outcome $\xi_i(\mathbf{x}, \mathbf{h})$, $i \in I$, using a two-step procedure that combines the FE model with TEs. TEs are machine learning models that combine predictions from multiple decision trees to capture nonlinear relationships and interactions among explanatory variables in 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.2 for a brief introduction.

In the first step, for each outcome i, we estimate an FE model analogous to (1), as described for ORDs and CIs in Section 4. From this model, we retain the estimated coefficients $\boldsymbol{\beta}^i$ and $\boldsymbol{\gamma}^i$ associated with the budget and socio-economic variables, respectively, that are statistically significant at the 10% level ($p < 0.1$), together with the estimated county and year fixed effects, $\alpha^i_{\tilde{c}}$ and $\kappa^i_{\tilde{t}+1}$. Let $\bar{J}^i$ and $\bar{S}^i_e$ denote the sets of statistically significant expenditure categories and socio-economic variables, respectively, for outcome i. We then represent the predicted outcomes as

$$\xi_i(\mathbf{x}, \mathbf{h}) = \sum_{j \in \bar{J}^i} \beta^i_j x_j + \sum_{s \in \bar{S}^i_e} \gamma^i_s h_s + \alpha^i_{\tilde{c}} + \kappa^i_{\tilde{t}+1} + \Phi^i(\mathbf{x}, \mathbf{h}), \qquad i \in I. \tag{4}$$

Here, the first four terms constitute the FE component of the prediction, while $\Phi^i(\mathbf{x}, \mathbf{h})$ represents the component learned by the TEs. Specifically, in the second step, we train a TEs to estimate $\Phi^i(\mathbf{x}, \mathbf{h})$. All explanatory variables included in the FE models are retained as inputs to the TEs, regardless of their statistical significance in Tables 3 and 4. For each observation k associated with outcome i, we first compute the residual from the FE component as

$$r^k_i = \xi^k_i - \left(\sum_{j \in \bar{J}^i} \beta^i_j x^k_j + \sum_{s \in \bar{S}^i_e} \gamma^i_s h^k_s + \alpha^i_{c(k)} + \kappa^i_{t(k)} \right),$$

where ξ^k_i denotes the observed value of outcome i for observation k, $c(k)$ and $t(k)$ denote the corresponding county and outcome year, respectively, and $(\mathbf{x}^k, \mathbf{h}^k)$ denotes the values of the explanatory variables for observation k in year $t(k) - 1$. We then use r^k_i as the dependent variable and $(\mathbf{x}^k, \mathbf{h}^k)$ as the input variables to train the TEs. The resulting function $\Phi^i(\mathbf{x}, \mathbf{h})$ therefore predicts the portion of outcome i not captured by the selected FE terms. The final prediction is obtained by adding this TEs prediction to the FE component, as specified above. Consequently, we refer to this model as FE + TEs.

The TE component $\Phi^i(\mathbf{x}, \mathbf{h})$ is trained using both the budget allocation variables $\mathbf{x}$ and the socio-economic variables $\mathbf{h}$. However, when solving the budget allocation problem, the socio-economic variables are not decision variables; rather, their values are known and fixed for the county and time period under consideration. Suppose that $\mathbf{h} = \hat{\mathbf{h}}$. Consider a node u in a tree of the TEs whose query involves a socio-economic variable h_s and is of the form $h_s \leq \upsilon$ versus $h_s > \upsilon$. Since $h_s = \hat{h}_s$ is fixed, only one of the two branches can be reached. Therefore, the split at node u can be removed by

replacing the node with the subtree corresponding to the branch satisfied by $\hat{h}_s$. Applying this procedure recursively to all such nodes removes all splits involving the socio-economic variables without changing the TEs predictions for a fixed $\hat{\mathbf{h}}$. Consequently, the resulting restricted TEs depends only on the budget allocation variables $\mathbf{x}$. For notational simplicity, we therefore denote $\Phi^i(\mathbf{x}, \hat{\mathbf{h}})$ by $\Phi^i(\mathbf{x})$ throughout the remainder of the paper.

5.1.1 Comparison of Response Trends Between the FE and FE+TEs Models

In this section, we examine how incorporating the TE component modifies the relationships between the input variables and the predicted outcomes relative to those identified by the FE models in Section 4. Our objective is to assess whether the response trends obtained from the combined FE+TEs model are consistent with the directions of the relationships identified by the FE models, while also examining how the TEs component introduces nonlinearities that are not represented by the linear FE model.

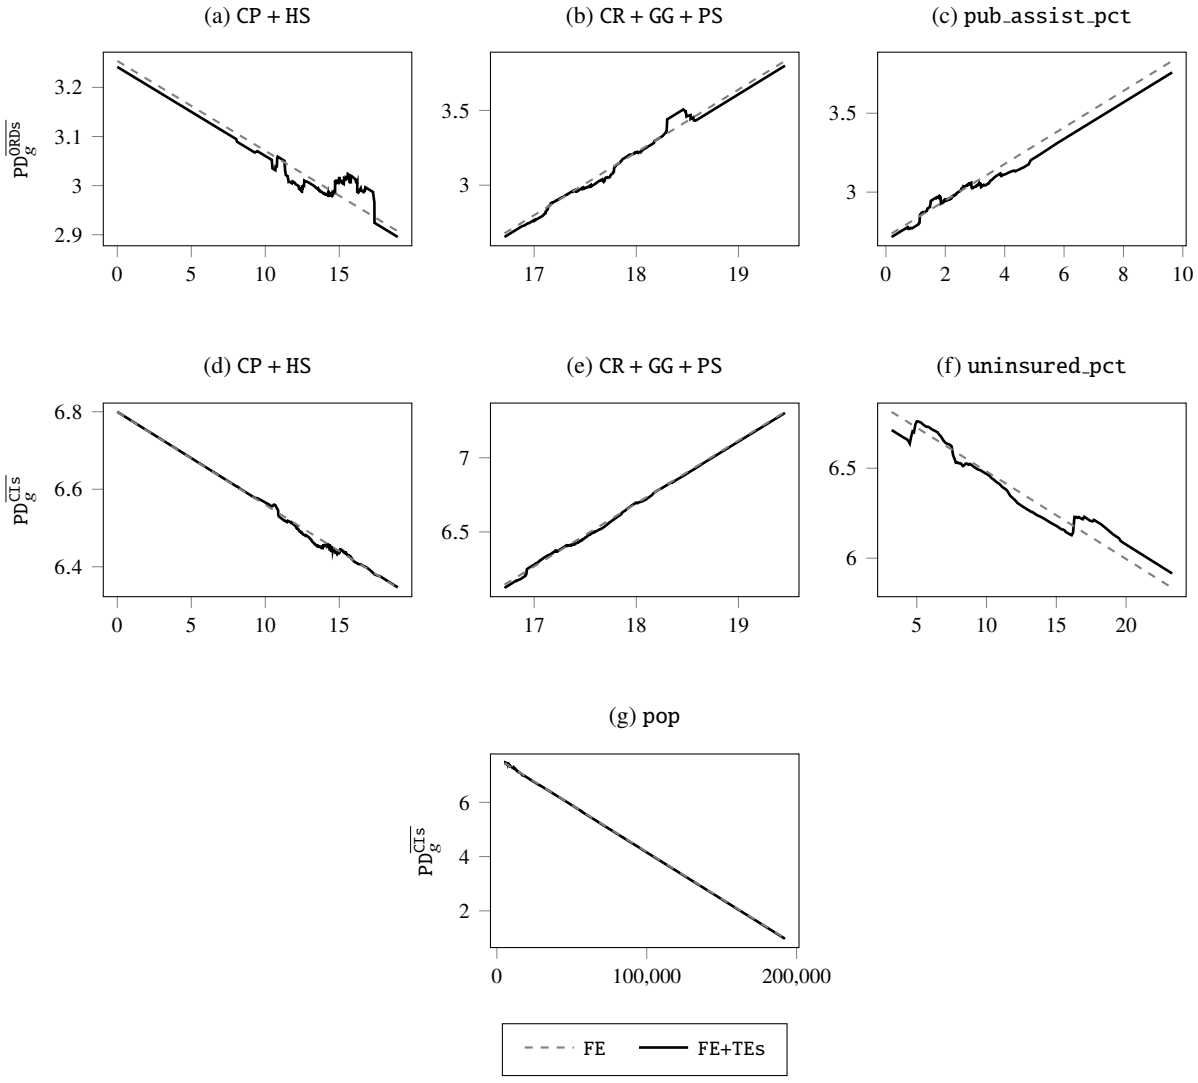


Figure 2: Comparison of response trends between the FE and FE+TEs models

To characterize the response of the FE+TEs model with respect to an individual explanatory variable, we use partial dependence functions (Friedman 2001, Ren et al. 2006). Let μ_g^q, $q = 1, \ldots, F$, denote the observed value of the explanatory variable g in the q^{th} observation where F is the total number of observations in the dataset. For the following presentation, we assume that the values of explanatory variable g are distinct across all observations in the dataset. The analysis can be easily extended to accommodate repeated values. For each observation k, let $\mathbf{v}^k = (\mathbf{x}^k, \mathbf{h}^k)$ denote the vector of all explanatory variables, and let $\mathbf{v}^k_{-g}$ denote the same vector excluding variable g. The partial dependence function for outcome i with respect to explanatory variable g, evaluated at μ_g^q, is defined as

$$\texttt{PD}_g^i(\mu_g^q) = \frac{1}{F}\sum_{k=1}^{F} \xi_i\left(\mu_g^q, \mathbf{v}^k_{-g}\right), \qquad q = 1, \ldots, F, \tag{5}$$

To evaluate the partial dependence function at μ_g^q, we set explanatory variable g to this value for all observations while retaining the observed values of the remaining explanatory variables. We then compute the corresponding FE+TEs predictions and average them across all observations. Repeating this procedure for each observed value of g produces the partial dependence curve, which summarizes how the predicted outcome varies with g, averaged over the observed values of the other explanatory variables.

Figure 2 presents these curves for the AECs and the socio-economic variables that are statistically significant in the FE models stated in (1). For comparison, we construct the corresponding FE curves using the same partial dependence procedure. Since the FE model is linear in the explanatory variables, this procedure produces a linear response with slope equal to the estimated FE coefficients in Tables 3 and 4. Comparing the FE and FE+TEs curves therefore allows us to assess whether the additional predictive structure captured by the TEs component preserves the direction of the relationship identified by the FE models and how it modifies the shape of that relationship.

For ORDs, the FE results reported in Table 3 identify a negative relationship with CP + HS, a positive relationship with CR + GG + PS, and a positive relationship with the percentage of individuals receiving public assistance. As shown in the top row of Figure 2, the corresponding FE+TEs partial dependence curves preserve these overall directions. Predicted ORDs generally decrease as CP + HS expenditures increase, whereas they increase as CR + GG + PS expenditures and public assistance increase. At the same time, the FE+TEs curves exhibit departures from the linear FE responses. These departures are particularly visible for CP + HS over the upper portions of its observed range, whereas the responses for CR + GG + PS and public assistance remain relatively closer to their corresponding linear FE trends.

A similar pattern is observed for CIs. The FE results in Table 4 identify a negative relationship between CIs and CP + HS and a positive relationship with CR + GG + PS. It also identifies negative relationships between CIs and both the percentage of the population without health insurance and total population. As shown in Figure 2, the FE+TEs response for CP + HS remains generally decreasing, while that for

CR + GG + PS remains generally increasing. Hence, for these two AECs, the FE+TEs model preserves the directions identified by the FE model while introducing nonlinear adjustments to the response curves. The uninsured percentage exhibits the most pronounced nonlinear departure from its corresponding linear FE trend. Although the FE model identifies an overall negative relationship, the FE+TEs partial dependence curve is non-monotonic over part of the observed range. For population, the FE+TEs response closely follows the decreasing linear FE trend, indicating little nonlinear departure from the relationship identified by the FE model.

Overall, the partial dependence analysis indicates that the FE+TEs model largely preserves the directions of the relationships identified by the FE models in Tables 3 and 4, while allowing their shapes and magnitudes to vary nonlinearly across the observed ranges of the explanatory variables. This comparison demonstrates how the TEs component complements the linear FE model by capturing additional predictive structure in the residual variation without substantially changing the overall response trends identified by the FE model. Section OS.3 demonstrates that FE+TEs achieves better out-of-sample predictive performance than standalone FE and TEs models.

5.1.2 Formulating Input-to-Output Relationships of TEs as MIPs

Given the values of socio-economic variables $\hat{\mathbf{h}}$ for county $\tilde{c}$ for year $\tilde{t}$, we are interested in modeling the relationship between $\mathbf{x}$ and $\mathbf{\Phi}^i$s over the set of possible budget allocations, and in particular, are interested in deriving an explicit MIP formulation for the set

$$O(\mathcal{T}, \mathbf{X}) = \left\{ (z, \mathbf{x}) \;\middle|\; z_i = \sum_{j \in \bar{J}^i} \beta^i_j x_j + \sum_{s \in \bar{S}^i_e} \gamma^i_s \hat{h}_s + \alpha^i_{\tilde{c}} + \kappa^i_{\tilde{t}+1} + \mathbf{\Phi}^i(\mathbf{x}),\ \forall i \in I,\ \mathbf{x} \in \mathbf{X} \right\}.$$

Next, we show that $O(\mathcal{T}, \mathbf{X})$ can be described as a mixed-integer linear set. This description extends the formulation in Kim et al. (2025), which itself built on the earlier work of Mišić (2020). Let $\mathcal{L}(\tau)$ denote the set of leaves of tree $\tau \in \mathcal{T}$ and L_j denote the collection of all split values involving expenditure category j that arise in any of the trees in $\mathcal{T}$. Specifically, for each $j \in J$, we assume without the loss of generality that $L_j = \{a_{j0}, a_{j1}, \ldots, a_{j(|L_j|-1)}\}$ where $0 = a_{j0} < a_{j1} < \cdots < a_{j(|L_j|-1)} = \infty$ and for each $v \in \{1, 2, \ldots, |L_j| - 2\}$ there exists a tree $\tau \in \mathcal{T}$ that includes the split condition $x_j \leq a_{jv}$. The elements of L_j divide the domain of expenditure category j into $|L_j|-1$ intervals of the form $(a_{j(v-1)}, a_{jv}]$ for $v \in \bar{L}_j := \{1, 2, \ldots, |L_j|-1\}$.

Consider a tree $\tau \in \mathcal{T}_i$. Each leaf of τ corresponds to a region in the input space $\mathbf{x}$ that is characterized by a collection of intervals on each of the expenditure categories. For each variable x_j and each leaf l of tree τ, we denote the set of these intervals by $K^j_{l\tau} \subseteq \bar{L}_j$. Finally, we let $c_{l\tau}$ be the output value of leaf l of tree τ and v_τ be the weight assigned to tree τ. The combination of sets $\mathcal{T}_i$, $K^j_{l\tau}$, L_j, and parameters $c_{l\tau}$ and v_τ, is sufficient to completely define and compute all values of all outcomes $\mathbf{\Phi}^i(\mathbf{x})$ for any $\mathbf{x}$.

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 $q_{l\tau}$ be a binary variable that takes the value 1 if leaf l of tree τ is selected and 0, otherwise. Additionally, let $\bar{x}_{jv}$ be a binary variable that takes the value 1 if x_j belongs to the interval $(a_{j(v-1)}, a_{jv}]$ and 0, otherwise. Finally, let proj $(\cdot)$ denote the projection operator. Building on Kim et al. (2025) and Mistry et al. (2021) and using the fact that expenditures are numerical, we write that $O(\mathcal{T}, \mathbf{X}) = \text{proj}_{(\mathbf{z},\mathbf{x})}\, \mathcal{P}(\mathcal{T}, \mathbf{X}) := \{(\mathbf{x}, \bar{\mathbf{x}}, \mathbf{z}, \mathbf{q}) \mid (6a) - (6h), \mathbf{x} \in \mathbf{X}\}$ whose feasible region contains the constraints

$$x_j \geq \sum_{v \in \bar{L}_j} a_{j(v-1)} \bar{x}_{jv} \qquad \forall j \in J \tag{6a}$$

$$x_j \leq \sum_{v \in \bar{L}_j} a_{jv} \bar{x}_{jv} \qquad \forall j \in J \tag{6b}$$

$$\sum_{v \in \bar{L}_j} \bar{x}_{jv} = 1 \qquad \forall j \in J \tag{6c}$$

$$\sum_{l \in \mathcal{L}(\tau): K^j_{l\tau} \subseteq \{v_1, \ldots, v_2\}} q_{l\tau} \leq \sum_{v=v_1}^{v_2} \bar{x}_{jv} \qquad \begin{array}{l} \forall \tau \in \mathcal{T}_i, \forall i \in I, \forall j \in J, \\ \forall v_1 \leq v_2 \in \bar{L}_j \end{array} \tag{6d}$$

$$\sum_{l \in \mathcal{L}(\tau)} q_{l\tau} = 1 \qquad \forall \tau \in \mathcal{T}_i, \forall i \in I \tag{6e}$$

$$z_i = \sum_{j \in \bar{J}^i} \beta^i_j x_j + \sum_{s \in \bar{S}^i_e} \gamma^i_s \hat{h}_s + \alpha^i_{\tilde{c}} + \kappa^i_{\tilde{t}+1} + \sum_{\tau \in \mathcal{T}_i} \sum_{l \in \mathcal{L}(\tau)} c_{l\tau} v_\tau q_{l\tau} \qquad \forall i \in I \tag{6f}$$

$$\mathbf{q}_{\cdot\tau} \in \{0,1\}^{|\mathcal{L}(\tau)|} \qquad \forall \tau \in \mathcal{T}_i, \forall i \in I \tag{6g}$$

$$\bar{\mathbf{x}}_j \in \{0,1\}^{|\bar{L}_j|} \qquad \forall j \in J. \tag{6h}$$

Constraints (6a)-(6c) ensure that, for each $j \in J$, an interval in $\bar{L}_j$ containing x_j is selected, whereas constraints (6d) and (6e) determine the leaf of each tree that is reached for a given $\bar{\mathbf{x}}_j$. Constraints (6f) aggregate the values of the selected leaves to compute $z_i = \xi_i(\mathbf{x}, \hat{\mathbf{h}})$. In the literature, the simplex constraint (6c) is often reformulated in incremental form to better leverage the branching capabilities of commercial solvers. Our models solve efficiently without this transformation.

Hence, the derivations in this section provide the following alternate formulation of (2):

$$\max_{\mathbf{x} \in \mathbf{X}} \min_{u \in \mathbf{U}} u(\bar{\xi}(\mathbf{x})) = \max_{(\mathbf{x}, \bar{\mathbf{x}}, \mathbf{z}, \mathbf{q}) \in \mathcal{P}(\mathcal{T}, \mathbf{X})} \min_{u \in \mathbf{U}} u(\mathbf{z}). \tag{7}$$

5.2 Robust Risk-Averse Utility Optimization Model

We now turn our attention to the inner problem of (7), $\min_{u \in \mathbf{U}} u(\mathbf{z})$, 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 (7).

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* $u(\mathbf{z}) = \sum_{i \in I} \sigma_i \bar{u}_i(z_i)$ *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* $\bar{u}_i(\cdot) \in \bar{\mathbf{U}}_i \ \ \forall i \in I$*.*

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 results in an additively separable utility function, a common approach in multi-attribute decision-making that combines evaluations of different outcomes according to their relative importance (Rezaei et al. 2022). This structure simplifies the modeling and computation by allowing the utility associated with each outcome $i \in I$ to be evaluated independently of the values of other outcomes $i' \in I$. 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 $\mathbb{R}^{|I|}$ that satisfy Assumption 1. Under this definition, the inner minimization problem in (7) reduces to

$$\min_{u \in \mathbf{U}} u(\mathbf{z}) = \sum_{i \in I} \sigma_i \bar{\pi}_i(\mathbf{z}) \quad \text{where} \quad \bar{\pi}_i(\mathbf{z}) = \min_{\bar{u}_i \in \bar{\mathbf{U}}_i} \bar{u}_i(z_i). \tag{8}$$

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. Given these partitions and Assumption 1, we approximate $\bar{\pi}_i(\mathbf{z})$ by solving the following linear program (Hu and Mehrotra 2015):

$$\begin{aligned}
\bar{\pi}_i(\mathbf{z}) = \min \quad & \tilde{\rho}_i z_i + \tilde{a}_i && && \text{(9a)}\\
s.t. \quad & && && \\
[\hat{\mathbf{y}}^1] \quad & \hat{\rho}_{i\iota+1} - \hat{\rho}_{i\iota} - \dot{\rho}_{i\iota}(p_{i\iota+1} - p_{i\iota}) \le 0, && \forall \iota = 0, \ldots, w_i - 1 && \text{(9b)}\\
[\hat{\mathbf{y}}^2] \quad & \hat{\rho}_{i\iota+1} - \hat{\rho}_{i\iota} - \dot{\rho}_{i\iota+1}(p_{i\iota+1} - p_{i\iota}) \ge 0, && \forall \iota = 0, \ldots, w_i - 2 && \text{(9c)}\\
[\hat{\mathbf{y}}^3] \quad & p_{i\iota}\tilde{\rho}_i + \tilde{a}_i \ge \hat{\rho}_{i\iota}, && \forall \iota = 0, \ldots, w_i && \text{(9d)}\\
[\hat{\mathbf{y}}^4], [\hat{\mathbf{y}}^5] \quad & \tilde{u}_i(p_{i\iota}, 0.41) \le \hat{\rho}_{i\iota} \le \tilde{u}_i(p_{i\iota}, 0.68), && \forall \iota = 0, \ldots, w_i && \text{(9e)}\\
[\hat{\mathbf{y}}^6], [\hat{\mathbf{y}}^7] \quad & \tilde{u}_i'(p_{i\iota}, 0.41) \le \dot{\rho}_{i\iota} \le \tilde{u}_i'(p_{i\iota}, 0.68), && \forall \iota = 0, \ldots, w_i - 1 && \text{(9f)}\\
& \hat{\rho}_{i\iota}, \tilde{a}_i, \tilde{\rho}_i, \dot{\rho}_{i\iota} \ge 0, && \forall \iota = 0, \ldots, w_i && \text{(9g)}
\end{aligned}$$

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

$$\tilde{u}_i(\omega, \lambda) = \frac{1}{1-\lambda}\omega^{1-\lambda}, \quad \lambda \in [0.41, 0.68] \tag{10}$$

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 $\tilde{\rho}_i$ and $\tilde{a}_i$ represent the slope and intercept of the linear segment of the piecewise-linear function at z_i. Variable $\hat{\rho}_{i\iota}$ captures the value of the function $\bar{\pi}_i$ at the point $p_{i\iota}$, whereas variable $\dot{\rho}_{i\iota}$ represents the slope of the line segment of the piecewise-linear function $\bar{\pi}_i$ over the interval $[p_{i\iota}, p_{i,\iota+1}]$. Constraints (9b) and (9c) ensure that $\bar{\pi}_i$ is piecewise-linear, concave, and non-decreasing. Constraints (9d) along with the objective function (9a) give the value of $\bar{\pi}_i(\mathbf{z})$. Constraints (9e) and (9f) 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 $\bar{u}_i^a(\cdot) = \tilde{u}_i(\cdot, 0.41)$ and $\bar{u}_i^b(\cdot) = \tilde{u}_i(\cdot, 0.68)$.

Next, we dualize (9) for each $i \in I$ to obtain an expression for $\bar{\pi}_i(\mathbf{z})$ as a maximization problem. To streamline 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 (9), 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

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

Using this notation, it is easily established that the dual of (9) is the optimization problem

$$\bar{\pi}_i(\mathbf{z}) = \max \left\{ \Psi_i(\hat{\mathbf{y}}_i^{1:7}) \;\middle|\; (\hat{\mathbf{y}}_i^1, \hat{\mathbf{y}}_i^2, \hat{\mathbf{y}}_i^3, \hat{\mathbf{y}}_i^4, \hat{\mathbf{y}}_i^5, \hat{\mathbf{y}}_i^6, \hat{\mathbf{y}}_i^7, z_i) \in \mathcal{R}_i \right\} \tag{11}$$

whose feasible region

$$\mathcal{R}_i = \left\{ (\hat{\mathbf{y}}_i^1, \hat{\mathbf{y}}_i^2, \hat{\mathbf{y}}_i^3, \hat{\mathbf{y}}_i^4, \hat{\mathbf{y}}_i^5, \hat{\mathbf{y}}_i^6, \hat{\mathbf{y}}_i^7, z_i) \;\middle|\; \text{(13a)} - \text{(13j)} \right\} \tag{12}$$

is defined through the constraints

$$-\hat{y}_{i0}^1 - \hat{y}_{i0}^2 - \hat{y}_{i0}^3 + \hat{y}_{i0}^4 + \hat{y}_{i0}^5 \leq 0 \tag{13a}$$

$$\hat{y}_{i(\iota-1)}^1 - \hat{y}_{i\iota}^1 + \hat{y}_{i(\iota-1)}^2 - \hat{y}_{i\iota}^2 - \hat{y}_{i\iota}^3 + \hat{y}_{i\iota}^4 + \hat{y}_{i\iota}^5 \leq 0 \qquad \forall \iota = 1, \ldots, w_i - 2 \tag{13b}$$

$$\hat{y}_{i(\iota-1)}^1 - \hat{y}_{i\iota}^1 + \hat{y}_{i(\iota-1)}^2 - \hat{y}_{i\iota}^3 + \hat{y}_{i\iota}^4 + \hat{y}_{i\iota}^5 \leq 0 \qquad \iota = w_i - 1 \tag{13c}$$

$$\hat{y}_{i(w_i-1)}^1 - \hat{y}_{iw_i}^3 + \hat{y}_{iw_i}^4 + \hat{y}_{iw_i}^5 \leq 0 \tag{13d}$$

$$-\hat{y}_{i0}^1(p_{i1} - p_{i0}) + \hat{y}_{i0}^6 + \hat{y}_{i0}^7 \leq 0 \tag{13e}$$

$$\begin{aligned} &-\hat{y}_{i\iota}^1(p_{i(\iota+1)} - p_{i\iota}) - \hat{y}_{i(\iota-1)}^2(p_{i\iota} - p_{i(\iota-1)}) \\ &\quad + \hat{y}_{i\iota}^6 + \hat{y}_{i\iota}^7 \leq 0 \qquad \forall \iota = 1, \ldots, w_i - 1 \end{aligned} \tag{13f}$$

$$\sum_{\iota=0}^{w_i} p_{i\iota} \hat{y}_{i\iota}^3 \leq z_i \tag{13g}$$

$$\sum_{\iota=0}^{w_i} \hat{y}_{i\iota}^3 \leq 1 \tag{13h}$$

$$\hat{\mathbf{y}}_i^1, \hat{\mathbf{y}}_i^5, \hat{\mathbf{y}}_i^7 \leq 0 \tag{13i}$$

$$\hat{\mathbf{y}}_i^2, \hat{\mathbf{y}}_i^3, \hat{\mathbf{y}}_i^4, \hat{\mathbf{y}}_i^6 \geq 0. \tag{13j}$$

5.3 MILP Formulation for (2) using TEs Trained to Predict Outcomes ξ

Combining (7), (8), and (11), we can reformulate (2) as the following MIP:

$$\begin{aligned} (\mathcal{F}) \qquad \max \quad & \sum_{i \in I} \sigma_i \Psi_i(\hat{\mathbf{y}}_i^{1:7}) \\ & (\hat{\mathbf{y}}_i^1, \hat{\mathbf{y}}_i^2, \hat{\mathbf{y}}_i^3, \hat{\mathbf{y}}_i^4, \hat{\mathbf{y}}_i^5, \hat{\mathbf{y}}_i^6, \hat{\mathbf{y}}_i^7, z_i) \in \mathcal{R}_i, \quad \forall i \in I \\ & (\mathbf{x}, \bar{\mathbf{x}}, \mathbf{z}, \mathbf{q}) \in \mathcal{P}(\mathcal{T}, \mathbf{X}). \end{aligned}$$

The expanded form of this model, after projecting out $\mathbf{z}$, can be found in Section OS.5 while Table OS.5 summarizes the key notation used in its derivation. 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 assess the performance of our budget allocation model $\mathcal{F}$ for 12 counties in southern WV for the year 2023: Boone, Cabell, Raleigh, Fayette, Kanawha, Lincoln, Logan, McDowell, Mercer, Mingo, Wayne, and Wyoming. These counties have been impacted the most in the state by the opioid epidemic, and they include communities historically dependent on the coal industry that have experienced significant economic disruption associated with its decline. The decline of the coal industry has contributed to economic distress and adverse health conditions in affected communities, factors that have been linked to opioid-related harms. At the same time, it has reduced local government revenues available to fund critical services and respond to the opioid crisis (National Association of Counties and Appalachian Regional Commission 2019, Love 2024, Young et al. 2023). The severity of the opioid crisis in this region is also evident in Figure OS.3b, which shows that southern WV experienced the highest average ORDs per 100,000 population in the state over the 2012-2022 period. All experiments in this section are performed on an Apple M2 processor with 8.00 GB of RAM, using Python 3.13.2. The optimization model is solved using GUROBI 11.0.0.

We incorporate a one-year lag between budget allocations and outcomes in the FE+TEs models. Specifically, when solving $\mathcal{F}$ for a given (county, year) pair, the model determines budget allocations for year t and evaluates the corresponding ORDs and CIs predicted for year $t + 1$. Accordingly, for our 2023 application, $\mathcal{F}$ determines budget allocations for 2022 and evaluates the corresponding predicted ORDs and CIs for 2023. For each county, we obtain the 2022 budget parameter b by summing the observed allocations per 100,000 population across all AECs. To preserve this temporal structure, we train two FE+TEs models, one for ORDs and one for CIs, using budget allocations and socio-economic characteristics from 2012-2021 and their corresponding outcomes from 2013-2022. The training features for the TEs component of each FE+TEs model include the log-transformed budget allocations per 100,000 population for AECs and all considered socio-economic variables. By retaining all explanatory variables, the TEs component can capture relationships between these variables and the outcomes that are not represented by the linear FE model.

The FE models estimated using outcomes till 2022 does not provide year fixed effects for 2023, which is needed to predict the 2023 outcomes using the FE+TEs model. Thus, following Bailey et al. (2025), we approximate this effect separately for each outcome using the three most recent years, calculated as the average of the estimated year fixed effects for 2020, 2021 and 2022. We use `scikit-learn` to train the TEs in FE+TEs models, assigning equal weights to all trees in the TE for each outcome and apply grid search for hyperparameter tuning. Further details of the hyperparameter tuning are provided in Section OS.3.

For each outcome $i \in I$, let θ_i^1 and θ_i^2 represent, respectively, the minimum and maximum predicted

values of outcome i attainable over the feasible budget allocations for the county under consideration. We determine these lower and upper bounds by solving the following optimization problems:

$$\theta_i^1 = \min \{z_i : (\mathbf{x}, \bar{\mathbf{x}}, \mathbf{z}, \mathbf{q}) \in \mathcal{P}(\mathcal{T}, \mathbf{X})\},$$
$$\theta_i^2 = \max \{z_i : (\mathbf{x}, \bar{\mathbf{x}}, \mathbf{z}, \mathbf{q}) \in \mathcal{P}(\mathcal{T}, \mathbf{X})\},$$

where $\mathcal{P}(\mathcal{T}, \mathbf{X})$ is the set defined by the constraints in (6). For each of the 12 counties considered in the 2023 analysis, we obtain $\theta_i^1 < \theta_i^2$ for both `ORDs` and `CIs`. We therefore partition the range $[\theta_i^1, \theta_i^2]$ into 20 equal-length intervals, yielding $\theta_i^1 = p_{i0} < p_{i1} < p_{i2} < \cdots < p_{iw_i} = \theta_i^2$, with $w_i = 20$.

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 $\tilde{\xi}_i(\mathbf{x}, \mathbf{h})$ for each $i \in I$ to reflect this preference structure. Specifically, we define the normalized outcomes as

$$\tilde{\xi}_i(\mathbf{x}, \mathbf{h}) = \frac{\eta\theta_i^2 - \xi_i(\mathbf{x}, \mathbf{h})}{\theta_i^2 - \theta_i^1} \qquad \forall i \in I. \tag{14}$$

The multiplier $\eta > 1$ in the numerator is included to ensure that $\tilde{\xi}_i(\mathbf{x})$ remains strictly positive, since the derivative of $\tilde{u}_i(\omega, \lambda)$, as defined in (10), is undefined at $\omega = 0$ for $\lambda \in [0.41, 0.68]$. In our computational experiments, we use $\eta = 1.2$.

In this analysis, we evaluate three pairs of coefficient values for $(\sigma_{\texttt{ORDs}}, \sigma_{\texttt{CIs}})$: $(1, 1)$, $(100, 1)$, and $(1, 100)$, each representing a different policy emphasis. We refer to $(1, 1)$ 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 define the reference budget allocations based on the fractions of the total budget allocated to each expenditure category in the previous year. Specifically, for each $j \in J$, we set $\hat{x}_j = \hat{\delta}_j b$, where $\hat{\delta}_j$ denotes the fraction of the total budget allocated to category j in the previous year. We set $\kappa_j^- = \kappa_j^+ = \delta b$ for all $j \in J$ and consider three values of $\delta \in \{0.02, 0.05, 0.10\}$. Thus, δ determines the maximum allowable deviation from the reference allocation for each expenditure category, expressed as a fraction of the total budget. For example, $\delta = 0.02$ allows the allocation to each category to deviate from its reference level by at most 2% of the total budget, whereas $\delta = 0.10$ permits a deviation of up to 10% of the total budget. Consequently, smaller values of δ restrict the extent to which $\mathcal{F}$ can reallocate the budget across `AECs`, while larger values provide greater flexibility in optimizing the outcomes.

Table OS.6 reports the mean computational time required to solve $\mathcal{F}$ across all combinations of *priority policy* and δ, with the maximum time shown in parentheses. As shown in the table, Gurobi solved all instances of $\mathcal{F}$ to optimality in less than 250 seconds. The remainder of this section is

structured as follows. Section 6.1 analyzes the reductions in predicted ORDs and CIs based on the budget allocations recommended by $\mathcal{F}$. Section 6.2 examines these budget allocations. Section 6.3 compares the recommended allocations with alternative budget allocation policies. Finally, Section 6.4 examines the sensitivity of the recommended allocations and predicted outcomes to changes in the available total budget. All results presented in this section are reported after converting the log-scale predictions from the FE+TEs models and the budget allocations from $\mathcal{F}$ to their original scales. Throughout this section, any references to the predictions and budget allocations obtained from these models are understood to be on their original scales.

6.1 Reduction in Predicted ORDs and CIs

In this section, we compare the predicted ORDs and CIs under the budget allocations identified by $\mathcal{F}$ with those under the observed budget allocations. We denote these predictions by $\widehat{\texttt{ORDs}}_{opt}$ and $\widehat{\texttt{CIs}}_{opt}$ for the optimized allocations and by $\widehat{\texttt{ORDs}}_{obs}$ and $\widehat{\texttt{CIs}}_{obs}$ for the observed allocations. Both sets of predictions are obtained using the trained FE+TEs models. To quantify the improvement, we compute the percentage reductions as $\zeta^{\widehat{\texttt{ORDs}}} = \frac{\widehat{\texttt{ORDs}}_{obs} - \widehat{\texttt{ORDs}}_{opt}}{\widehat{\texttt{ORDs}}_{obs}} \times 100$ and $\zeta^{\widehat{\texttt{CIs}}} = \frac{\widehat{\texttt{CIs}}_{obs} - \widehat{\texttt{CIs}}_{opt}}{\widehat{\texttt{CIs}}_{obs}} \times 100$. We begin with a detailed analysis of Cabell County as an illustrative case before presenting results for all 12 counties. We select Cabell because it has experienced the highest per-capita levels of ORDs, relatively high per-capita levels of CIs and a comparatively low per-capita total budget, providing a useful setting for examining how $\mathcal{F}$ reallocates a fixed budget to reduce the predicted levels of both outcomes.

Figure 3 presents the values of $\zeta^{\widehat{\texttt{ORDs}}}$ and $\zeta^{\widehat{\texttt{CIs}}}$ for Cabell County in 2023 under each of the three *priority policies* (BPP, OPP, and CPP) and for each of the three values of δ. Under BPP, $\mathcal{F}$ yields reductions in both predicted outcomes for $\delta = 0.02$ and $\delta = 0.05$. Specifically, $\zeta^{\widehat{\texttt{ORDs}}}$ is 2.04% and 1.41%, respectively, while $\zeta^{\widehat{\texttt{CIs}}}$ is 5.95% and 11.48%. When δ increases to 0.10, the reduction in predicted CIs increases further to 18.50%, whereas $\zeta^{\widehat{\texttt{ORDs}}} = -0.88\%$, indicating a slight increase in predicted ORDs relative to the observed budget allocation.

Increasing δ generally allows the model greater flexibility to improve the prioritized outcome, as reflected by the non-decreasing reductions in ORDs under OPP and in CIs under CPP. Furthermore, the results under OPP and CPP also provide an interesting connection to our motivation for considering both ORDs and CIs in the framework. Under OPP, $\zeta^{\widehat{\texttt{ORDs}}}$ increases from 2.36% at $\delta = 0.02$ to 6.43% at $\delta = 0.05$ and 8.88% at $\delta = 0.10$. These reductions in predicted ORDs are accompanied by reductions in predicted CIs, with $\zeta^{\widehat{\texttt{CIs}}}$ increasing from 2.29% to 6.60% over the same range of δ. This pattern is consistent with prior research documenting links between opioid misuse and criminal activity, as well as evidence that substance-use and opioid-related interventions can have spillover effects on crime (Bondurant et al. 2018, Dave et al. 2021, Sim 2023). The converse, however, need not hold because CIs capture criminal activity

arising from a broader set of factors, many of which are unrelated to opioid misuse. This distinction is reflected under CPP, which produces the largest reductions in predicted CIs, with $\zeta^{\widehat{\texttt{CIs}}}$ increasing from 16.22% at $\delta = 0.02$ to 20.83% at $\delta = 0.10$, while $\zeta^{\widehat{\texttt{ORDs}}}$ remains negative. Thus, allocations that are favorable for reducing predicted CIs are not necessarily favorable for reducing predicted ORDs. These results reinforce the importance of considering both outcomes jointly when evaluating alternative budget allocations.

Overall, the results demonstrate how the choice of *priority policy* affects the allocation recommendations generated by $\mathcal{F}$. Assigning greater priority to ORDs through OPP produces larger reductions in predicted ORDs than BPP for $\delta = 0.05$ and $\delta = 0.10$, but smaller reductions in predicted CIs. Conversely, prioritizing CIs through CPP produces substantially larger reductions in predicted CIs, but at the expense of higher predicted ORDs. The results therefore illustrate the trade-off between the two outcomes. 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.

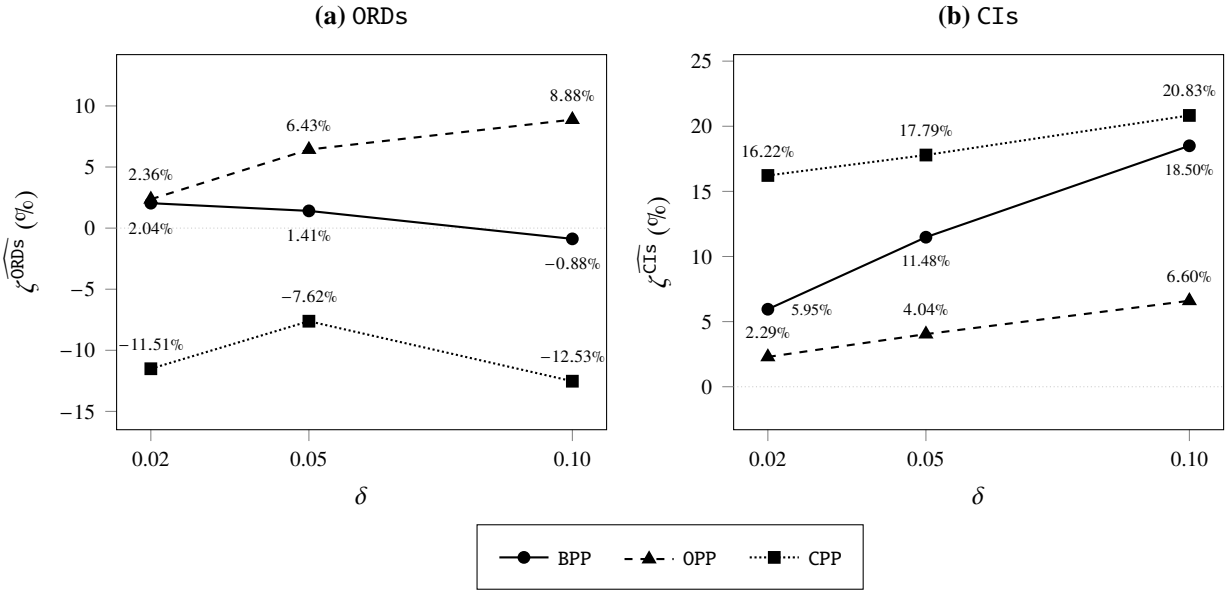


Figure 3: Percentage reductions in predicted ORDs and CIs under the three *priority policies* and δ values for Cabell County in 2023.

Next, we examine whether the patterns observed for Cabell County extend to the other 11 counties in

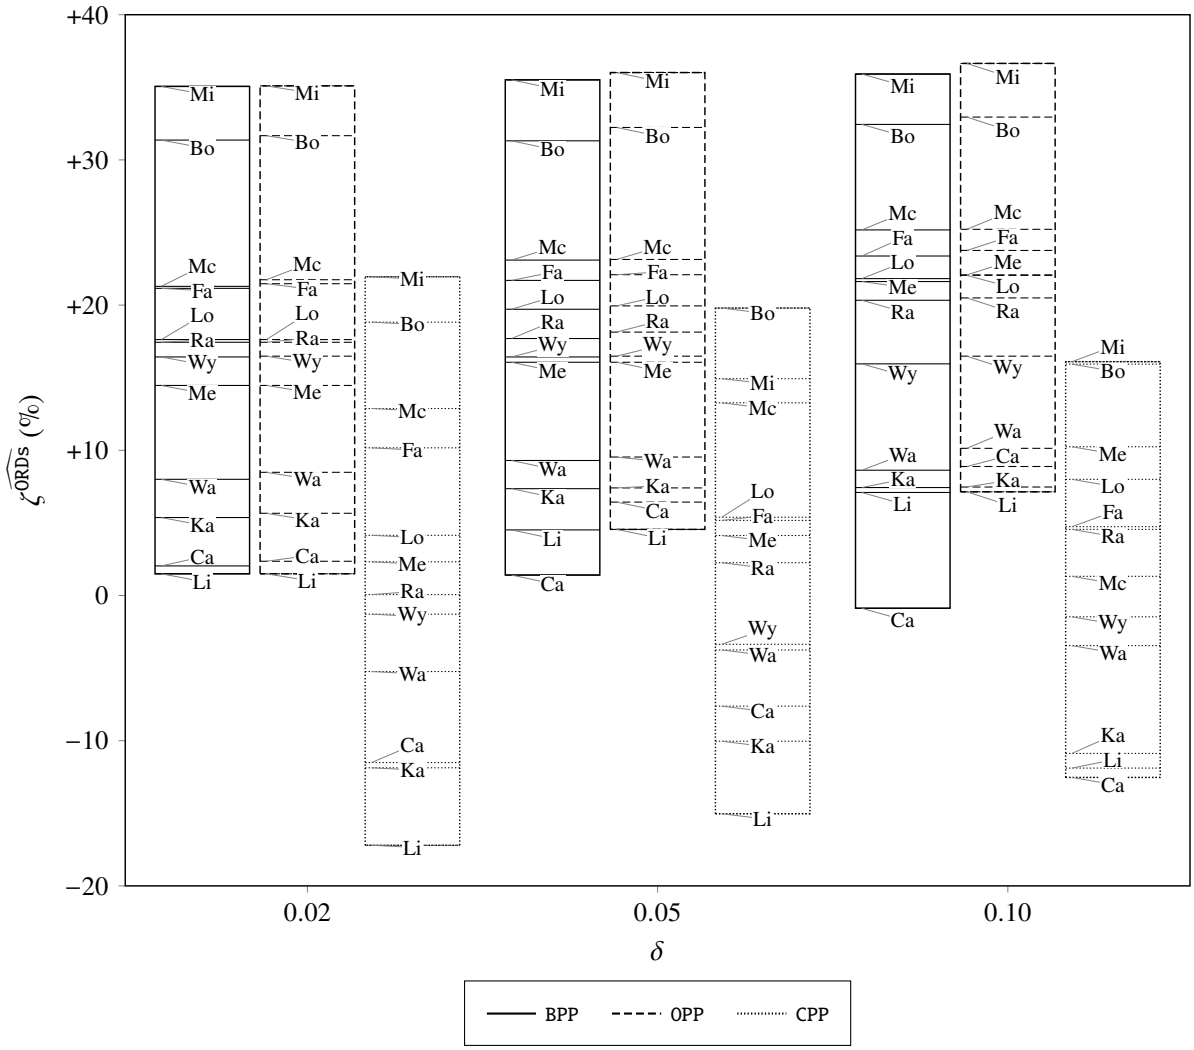


Figure 4: Percentage reductions in predicted `ORDs` under the three *priority policies* and δ values for 12 southern counties in `WV`. County abbreviations: Bo = Boone, Ca = Cabell, Fa = Fayette, Ka = Kanawha, Li = Lincoln, Lo = Logan, Mc = McDowell, Me = Mercer, Mi = Mingo, Ra = Raleigh, Wa = Wayne, and Wy = Wyoming.

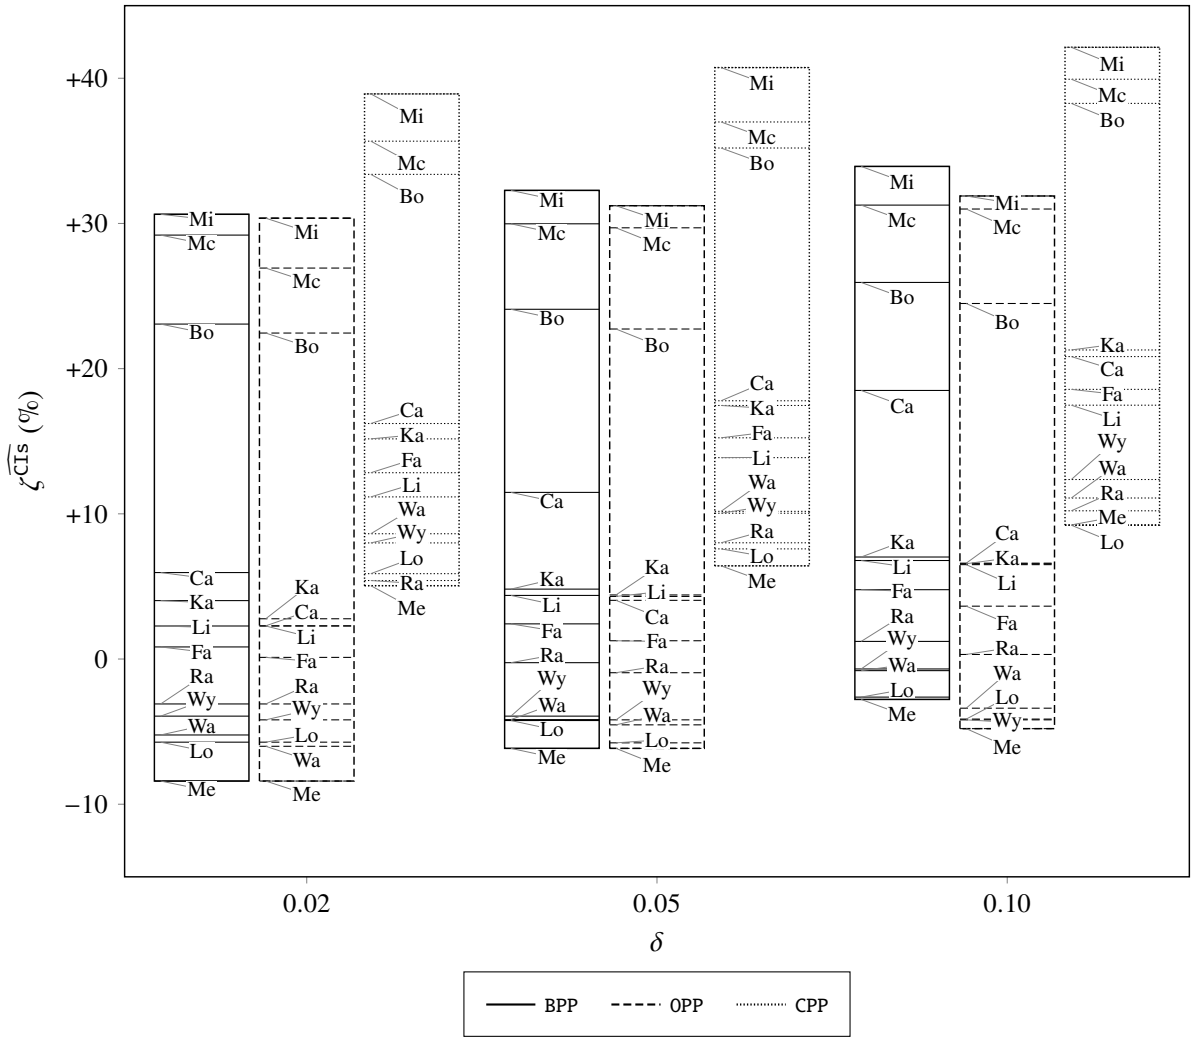


Figure 5: Percentage reductions in predicted `CIs` under the three *priority policies* and δ values for 12 southern counties in `WV`. County abbreviations: Bo = Boone, Ca = Cabell, Fa = Fayette, Ka = Kanawha, Li = Lincoln, Lo = Logan, Mc = McDowell, Me = Mercer, Mi = Mingo, Ra = Raleigh, Wa = Wayne, and Wy = Wyoming.

southern WV. Figures 4 and 5 summarize the results for all 12 counties. Overall, $\mathcal{F}$ identifies budget allocations associated with reductions in predicted ORDs and CIs for most counties, although the magnitude of these reductions varies considerably across counties and *priority policies*. Under OPP, $\zeta^{\widehat{\texttt{ORDs}}}$ is positive for all 12 counties and generally increases as δ increases. Similarly, under CPP, $\zeta^{\widehat{\texttt{CIs}}}$ is positive for all counties and generally increases with δ. The results also reveal a trade-off between the two outcomes across many counties. Relative to BPP, OPP generally produces larger reductions in predicted ORDs but smaller reductions in predicted CIs, whereas CPP generally produces larger reductions in predicted CIs but smaller reductions in predicted ORDs. Thus, prioritizing one outcome can reduce, and in some cases reverse, the improvement in the other. At the same time, the extent of these trade-offs varies across counties, highlighting the heterogeneity in the allocation recommendations and the importance of accounting for county-specific conditions.

6.2 Changes in Budget Allocations

We now examine the budget allocations underlying the predicted ORDs and CIs reported in Section 6.1. We define the *expenditure share* of a category as the fraction of a county's total budget allocated to that expenditure category. We express the allocations as *expenditure shares* of each county's total 2022 budget to facilitate comparisons across expenditure categories and counties.

We first examine Cabell County in detail. Table 5 reports the observed and optimized *expenditure shares* for AECs under each combination of *priority policy* and δ. Each entry reports the optimized 2022 *expenditure shares*, while the observed 2022 *expenditure shares* are reported at the bottom of the table. The observed allocation in Cabell County is highly concentrated in CR + GG + PS, which accounts for 99.60% of the total budget, compared with 0.40% for CP + HS and 0% for SS. Across all three *priority policies*, $\mathcal{F}$ generally reallocates a portion of Cabell County's budget from CR + GG + PS toward CP + HS, with larger values of δ allowing greater reallocations. The magnitude and composition of these adjustments vary across the *priority policies*; in particular, CPP produces the largest shift toward CP + HS and also allocates a positive share to SS. These differences in the optimized budget allocations across the *priority policies* and values of δ correspond to the differences in predicted ORDs and CIs reported in Figure 3.

More importantly, the recommended allocations are not simply determined by the signs or statistical significance of the coefficients in the FE models. The FE results indicate a negative association between CP + HS expenditures and both ORDs and CIs, whereas CR + GG + PS expenditures exhibit a positive association with both outcomes. The coefficients for SS are not statistically significant for either outcome. Across all three *priority policies*, however, the optimized allocations exhibit patterns that cannot be inferred from these linear relationships alone. For example, under BPP with $\delta = 0.10$, the share allocated

to CR + GG + PS decreases from 99.60% to 89.68%, nearly exhausting the allowable 10% reduction under the RBA constraint, while the share allocated to CP + HS increases only from 0.40% to 4.91% and 2.61% is allocated to SS. Under OPP, increasing δ from 0.05 to 0.10 provides additional allocation flexibility, yet the optimized share of CP + HS remains unchanged at 5.04%, while SS continues to receive no allocation. Similarly, under CPP, SS receives a positive allocation for all three values of δ, reaching 1.12% at $\delta = 0.10$, despite its statistically insignificant FE coefficients. The resulting recommendations therefore reflect the combined effect of the FE+TEs models and $\mathcal{F}$, rather than a direct translation of individual FE coefficients into budget changes.

We next examine whether the allocation patterns observed for Cabell County extend across all remaining 11 counties in southern WV. Figure 6 compares the mean observed and optimized *expenditure shares* across the counties for each expenditure category, *priority policy*, and value of δ. The observed allocations are highly concentrated in CR + GG + PS, which accounts for an average of 98.54% of the total budget, compared with 1.27% for CP + HS and 0.19% for SS (see Table OS.2). Across all three *priority policies*, $\mathcal{F}$ generally shifts *expenditure share* away from CR + GG + PS and toward CP + HS, with larger values of δ allowing greater adjustments. The allocation to SS, however, exhibits a different pattern, remaining at or near zero under BPP and OPP while receiving a positive mean share under CPP for all three values of δ. These results show that the mean allocation patterns are broadly consistent with the directions of the FE coefficients but also reflect the role of FE+TEs models and $\mathcal{F}$. The detailed county-level results are stated in Table OS.7.

Priority policy	δ	CP+HS (%)	CR+GG+PS (%)	SS (%)	ρ_c
BPP	0.02	2.32	97.68	0.00	1.000
	0.05	3.31	95.12	0.00	0.984
	0.10	4.91	89.68	2.61	0.972
OPP	0.02	0.70	98.39	0.00	0.991
	0.05	5.04	94.68	0.00	0.997
	0.10	5.04	89.68	0.00	0.947
CPP	0.02	1.73	97.68	0.59	1.000
	0.05	3.10	95.75	0.59	0.994
	0.10	9.16	89.68	1.12	1.000
Observed *expenditure shares*		0.40	99.60	0.00	
Total budget		$29,288,375.50			

Table 5: Observed and optimized 2022 *expenditure shares* for Cabell County under the three *priority policies* and δ values. ρ_c is the fraction of available budget spent in optimized allocation.

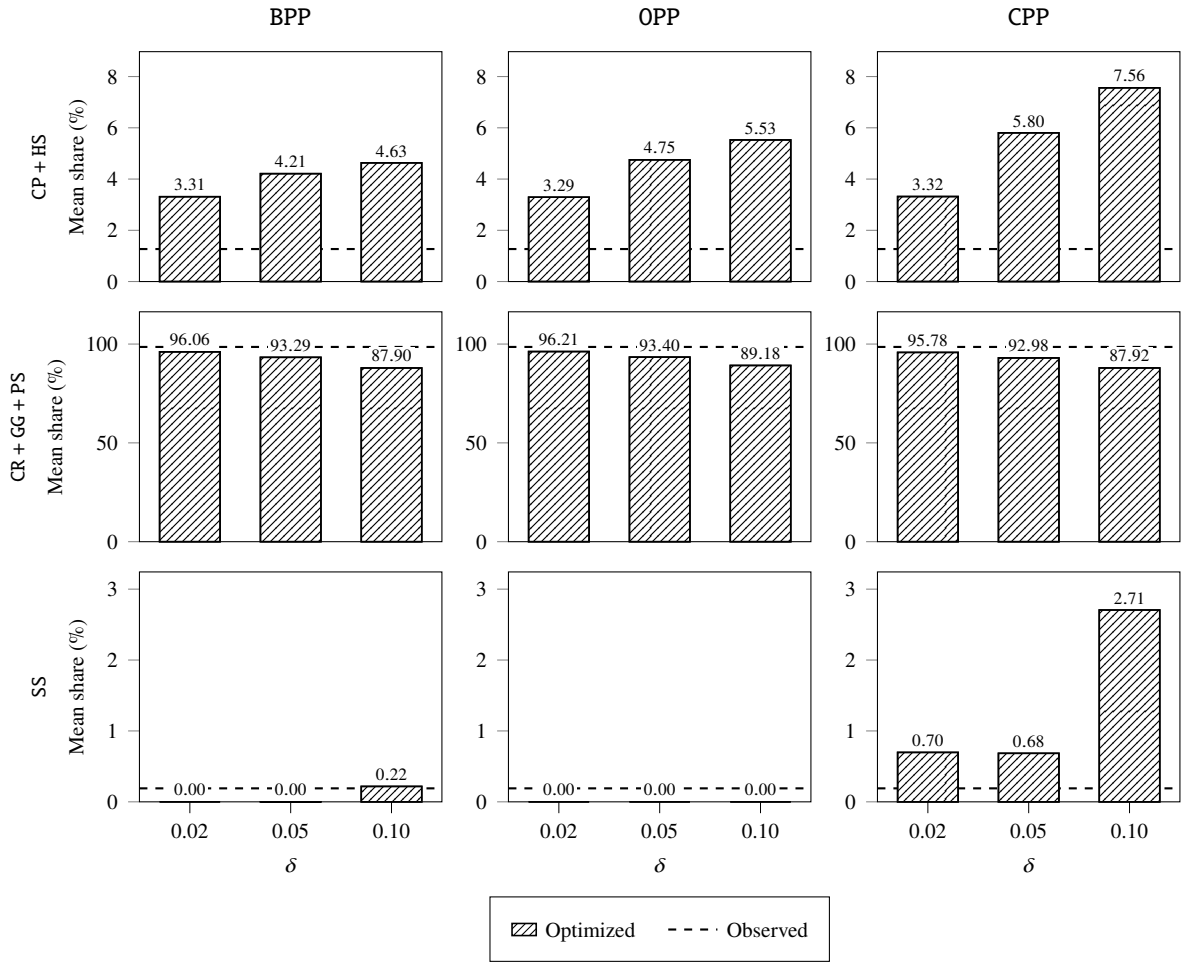


Figure 6: Mean observed and optimized 2022 *expenditure shares* across 12 southern counties in WV under the three *priority policies* and δ values.

6.3 Comparison with Alternative Budget Allocation Policies

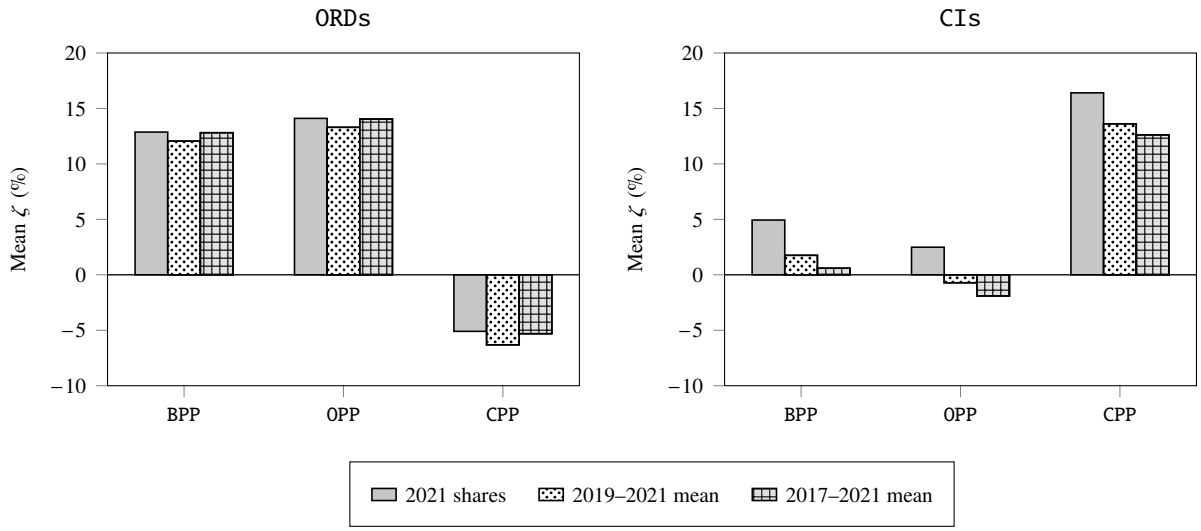


Figure 7: Mean percentage differences in predicted ORDs and CIs between the alternative allocation policies and the allocations identified by $\mathcal{F}$ under $\delta = 0.10$, averaged across 12 southern counties in WV.

To further evaluate the allocations identified by $\mathcal{F}$, we compare their predicted outcomes with those associated with three alternative budget allocation policies based on historical expenditure patterns. For each county, we consider alternative 2022 budget allocations based on (i) *expenditure shares* observed in 2021, (ii) mean *expenditure shares* over 2019–2021, and (iii) mean *expenditure shares* over 2017–2021. These three policies therefore provide alternative 2022 budget allocations reflecting the most recent, three-year average, and five-year average historical spending patterns, respectively. Because these alternative policies are not subject to the RBA constraints, we compare them with the allocations identified by $\mathcal{F}$ under $\delta = 0.10$, the most flexible allocation setting considered in our analysis.

For each alternative policy, the corresponding 2022 budget allocations are used as inputs to the trained FE+TEs models to obtain the predicted outcomes for 2023. For outcome $i \in I$, county c, and alternative allocation policy a, we calculate $\zeta_{ica} = \frac{\xi_{ica}^{\text{alt}} - \xi_{ic}^{\text{opt}}}{\xi_{ica}^{\text{alt}}} \times 100$. Here ξ_{ica}^{alt} denotes the predicted value of outcome i under alternative allocation policy a , and ξ_{ic}^{opt} denotes the predicted value under the allocation identified by $\mathcal{F}$. Figure 7 reports the mean of ζ_{ica} across the 12 counties for each alternative allocation policy under BPP, OPP, and CPP. A positive value indicates that the predicted outcome under the alternative policy is higher than under the optimized allocation, whereas a negative value indicates that it is lower.

The results show a consistent pattern across the three alternative allocation policies. Under both BPP and OPP, the optimized allocations yield substantially lower predicted ORDs, with mean reductions of at least 12% and 13%, respectively. In comparison, the differences in predicted CIs are relatively modest: depending on the alternative policy, the optimized allocations result in either a small reduction or a small increase in predicted CIs. Under CPP, the pattern is reversed. The optimized allocations produce

substantially lower predicted CIs than all three alternative allocation policies, with mean reductions of at-least 12%. At the same time, predicted ORDs under the optimized allocations are approximately 5% higher than under the alternative policies. This result is consistent with the trade-off identified in the preceding analysis: allocations that place greater priority on reducing predicted CIs can lead to less favorable predicted ORDs.

6.4 Sensitivity to the Available Budget

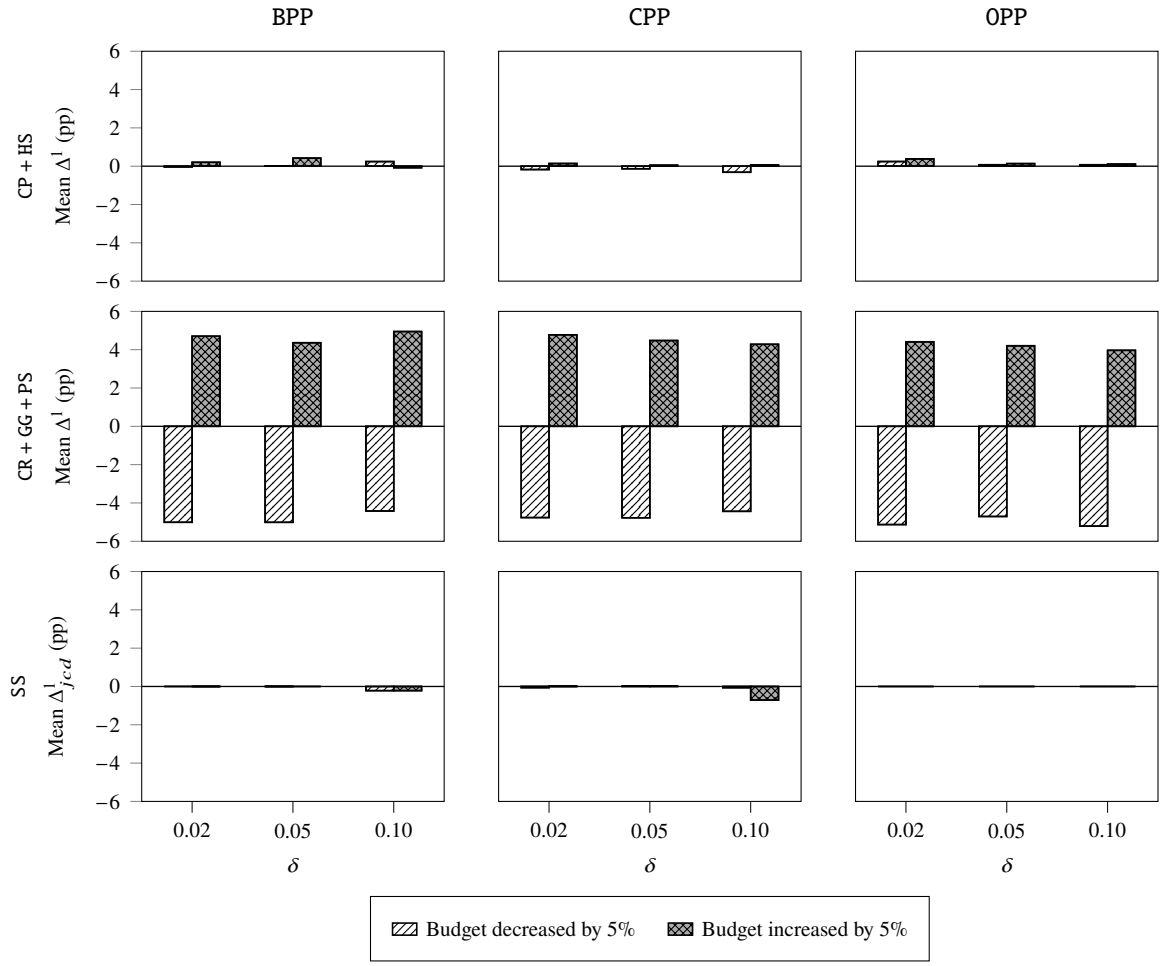


Figure 8: Mean Δ^1 across 12 southern counties in WV under alternative budget levels.

We examine the sensitivity of the budget allocations recommended by $\mathcal{F}$ and the corresponding predicted outcomes to changes in the total budget available to each county. Specifically, for each of the 12 counties, we consider two alternative budget levels, obtained by decreasing and increasing the available 2022 budget by 5%. We resolve $\mathcal{F}$ at each budget level for all three *priority policies* and each value of δ. To measure the resulting changes in optimized budget allocations, for expenditure category j, county c, and budget level $d \in \{-5\%, +5\%\}$, we calculate $\Delta^1_{jcd} = \frac{\hat{x}_{jcd} - x^*_{jc}}{b^*_c} \times 100$, where $\hat{x}_{jcd}$ denotes the optimized allocation to expenditure category j under budget level d, x^*_{jc} denotes the corresponding

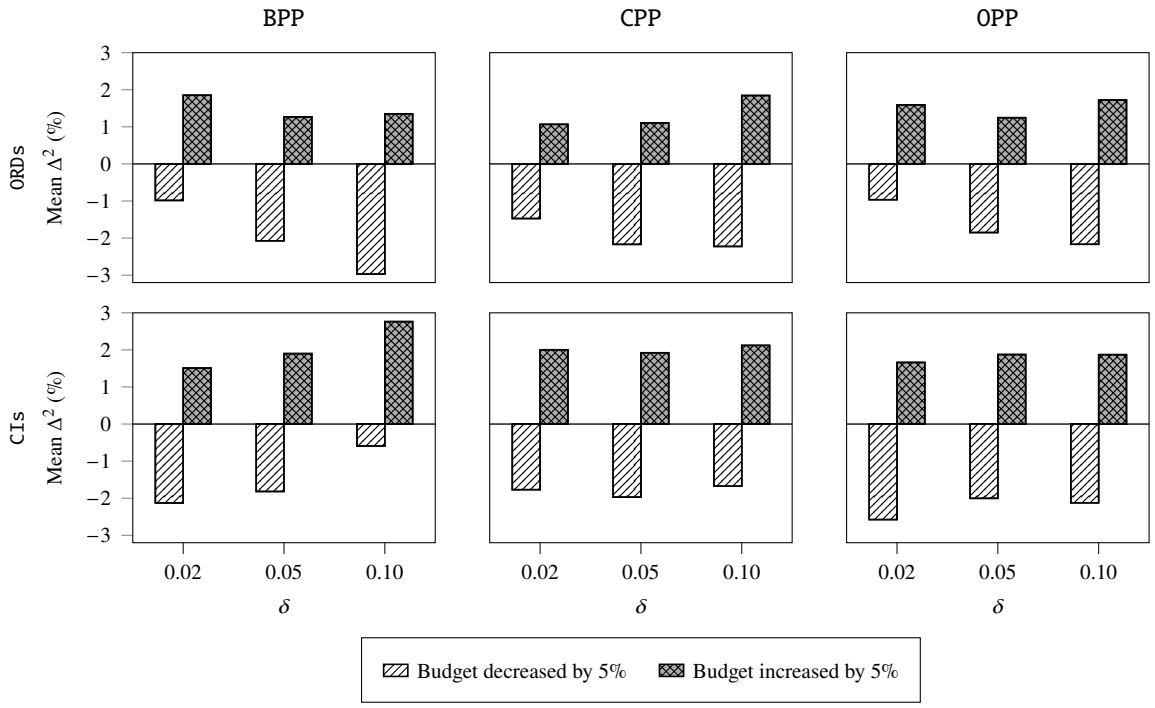


Figure 9: Mean Δ^2 across 12 southern counties in WV under alternative budget levels.

optimized allocation under the unchanged budget, and b_c^* is the unchanged 2022 total budget of county c. Thus, Δ^1_{jcd} measures the change in the optimized allocation to category j as a percentage of the unchanged 2022 budget. Similarly, to measure changes in predicted outcomes, for each outcome $i \in I$, we calculate $\Delta^2_{icd} = \frac{\xi^{\text{opt}}_{icd} - \xi^*_{ic}}{\xi^*_{ic}} \times 100$, where ξ^{opt}_{icd} and ξ^*_{ic} denote the predicted values of outcome i for county c under budget level d and the unchanged budget, respectively. A positive (negative) value of Δ^2_{icd} indicates an increase (decrease) in the predicted outcome relative to the unchanged-budget case.

Figures 8 and 9 report mean Δ^1_{jcd} and Δ^2_{icd} across the 12 counties, respectively, for each outcome, *priority policy*, δ and budget level d. In these figures, a clear pattern emerges across the three *priority policies*. Changes in the available budget are reflected primarily in CR + GG + PS, whereas the optimized allocations to CP + HS and SS remain comparatively stable. When the available budget is reduced by 5%, the allocation to CR + GG + PS decreases substantially across all *priority policies* and values of δ. Conversely, when the available budget is increased by 5%, most of the additional resources are allocated to CR + GG + PS. The corresponding outcome results exhibit a counterintuitive pattern: increasing the available budget generally results in higher predicted ORDs and CIs, whereas reducing the available budget generally results in lower predicted levels of both outcomes.

This pattern can be explained in part by the structure of the RBA constraints. For the 2022 allocation decisions, the reference allocation for each expenditure category is determined using its share of the county's total budget in 2021. Consequently, when the total available budget changes, the reference allocation and the corresponding bounds for each expenditure category in the RBA constraints change

with it. Because CR + GG + PS accounts for the vast majority of the 2021 budgets, an increase in the total available budget produces a substantially larger increase in its reference allocation than in those of CP+HS and SS. The RBA constraints limit the extent to which $\mathcal{F}$ can deviate from these reference allocations, resulting in most of the additional budget being allocated to CR + GG + PS. The reverse occurs when the available budget is reduced. Combined with the relationships represented by the FE+TEs models, this allocation mechanism explains why a larger available budget can result in higher predicted ORDs and CIs. Importantly, this result should not be interpreted as indicating that greater county resources inherently worsen opioid-related outcomes. Rather, it shows that the effect of additional resources depends on how those resources can be distributed across expenditure categories.

From a practical perspective, this behavior is also consistent with the incremental nature of public-sector budgeting. In incremental budgeting, existing expenditure levels provide the base for subsequent budget decisions, with attention focused primarily on increases or decreases relative to that base. This approach has also been documented in local-government expenditure decisions (Kelly 2005, Good 2011). The RBA constraints reflect this feature of the budget-setting process by using the 2021 *expenditure shares* as the reference for the 2022 allocation decisions and restricting the magnitude of deviations from those shares. Consequently, an increase in the available budget does not necessarily imply that the additional resources can be redirected toward the AECs associated with the greatest improvements in the considered outcomes. This highlights an important practical implication of the framework: the benefits of additional resources may depend not only on the amount of resources available but also on the flexibility policymakers have to reprioritize those resources across AECs.

7 Conclusion

In this paper, we develop a county-level decision-support framework that integrates FE models, 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 employ FE models to estimate the effects of AECs and socio-economic variables on ORDs and CIs, while accounting for county and year fixed effects. We then combine the linear FE models with TEs that capture remaining relationships between outcomes and the explanatory variables. Subsequently, for a given county, we embed the trained FE+TEs models into a MIP formulation that generates budget allocations across AECs 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 FE models identify statistically significant relationships between budget allocations, socio-economic variables, and both ORDs and CIs. In our computational experiments, we apply $\mathcal{F}$ to 12 southern counties in WV that have been most impacted by the opioid crisis, using AECs: CP + HS, CR + GG + PS, and SS, and targeting two outcomes: ORDs and CIs. 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). We consider three values of $\delta \in \{0.02, 0.05, 0.10\}$, where δ determines the maximum allowable deviation from the reference allocation for each expenditure category, with larger values providing greater flexibility to reallocate expenditures across categories. Our computational results indicate that $\mathcal{F}$ can effectively recommend budget allocations across AECs to improve the predicted outcomes. Our analysis also reveals a trade-off between the two outcomes: greater improvements in one may lead to lower improvements or deterioration in the other. Finally, our sensitivity analysis of the total available budget reveals that changes in budget availability are primarily absorbed by CR + GG + PS in the optimal allocations generated by $\mathcal{F}$, while allocations to the other expenditure categories remain relatively stable. Interestingly, increasing the available budget does not necessarily improve the predicted outcomes, highlighting that the benefits of additional resources depend on the flexibility to reallocate them across AECs.

This work has some limitations. First, the current analysis focuses on the relationship between budget allocations and outcomes observed one year later; therefore, it may not capture longer-term effects of county expenditures. For example, investments in CR, such as youth development, arts, and community activities, may generate benefits over a longer time horizon. We examined the FE model with budget allocations lagged by more than one year; however, these longer lags did not yield statistically significant relationships in our data. Second, although $\mathcal{F}$ incorporates an available budget, allocation limits, and alternative policymaker priorities, it does not capture all institutional and operational considerations of actual county budget-setting processes, such as statutory spending obligations, commitments from prior fiscal years, and restrictions on transferring resources across AECs. The current formulation also does not explicitly incorporate fairness or equity considerations in the distribution of public resources or their resulting outcomes. Accordingly, the allocations generated by $\mathcal{F}$ should be viewed as decision-support recommendations rather than direct prescriptions for county budgets.

Building on these limitations, we identify several directions for future research. First, the framework could be extended to incorporate longer-term effects of budget allocations as additional longitudinal data become available, allowing investments whose benefits may emerge over multiple years to be better represented. Second, additional institutional constraints and fairness or equity considerations could be incorporated into $\mathcal{F}$ to more closely reflect actual county budget-setting processes. Third, incorporating

additional outcome metrics, such as the share of individuals in recovery, public confidence in law enforcement, and economic growth, could provide a more holistic view of how budget allocations help in managing the opioid crisis. Fourth, as `AECs` contain multiple subcategories, the framework can accommodate budget allocation at a more granular level. One practical extension is a two-step approach in which the model first identifies which major budget categories should receive increased or decreased funding—and by what amount—and then determines how these changes should be distributed among the subcategories within each major category. Finally, the current approach assumes that policymakers are risk-averse in their decision-making. Future work could relax this assumption to capture more nuanced, potentially risk-seeking preferences.

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.

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.

Bailey, N., Z. Hochman, Y. Mao, M. Silvapulle, and P. Silvapulle (2025). Impact of climate change on agriculture in australia: an interactive fixed effects model approach. *Applied Economics 57*(43), 6901–6914.

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.

Bondurant, S. R., J. M. Lindo, and I. D. Swensen (2018). Substance abuse treatment centers and local crime. *Journal of Urban Economics 104*, 124–133.

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.

Centers for Disease Control and Prevention (2025). Opioid overdose data. `https://www.cdc.gov/nchs/`.

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.

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.

Dong, Q., D. Kline, and S. A. Hepler (2023). A bayesian spatio-temporal model to optimize allocation of buprenorphine in north carolina. *Statistics and Public Policy 10*(1), 2218448.

Drug Enforcement Administration (2025). 2025 national drug threat assessment summary. 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.

Friedman, J. H. (2001). Greedy function approximation: a gradient boosting machine. *Annals of Statistics*, 1189–1232.

Garnett, M. F. and A. M. Miniño (2025). Changes in drug overdose mortality and selected drug type by state: United states, 2022 to 2023. NCHS Health E-Stat 101.

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.

Good, D. A. (2011). Still budgeting by muddling through: Why disjointed incrementalism lasts. *Policy and Society 30*(1), 41–51.

Hair, J. F., B. J. Babin, R. E. Anderson, and W. C. Black (2022). *Multivariate Data Analysis* (8th ed.). Cengage Learning.

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

Hedegaard, H., M. Warner, and A. M. Miniño (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.

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

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.

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.

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

Kelly, J. M. (2005). A century of public budgeting reform: The "key" question. *Administration & Society 37*(1), 89–109.

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.

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.

Love, H. (2024). Catalyzing a 'just transition' from coal to clean energy in west virginia's coalfields. *Brookings*.

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.

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), 202–207.

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.

Mizan, T. and S. Taghipour (2022). Medical resource allocation planning by integrating machine learning and optimization models. *Artificial Intelligence in Medicine 134*, 102430.

National Association of Counties and Appalachian Regional Commission (2019). Opioids in appalachia: The role of counties in reversing a regional epidemic. Technical report, National Association of Counties and Appalachian Regional Commission.

National Institute on Minority Health and Health Disparities (2026). HD*Pulse: An ecosystem of health disparities and specific populations health resources. Created October 19, 2026.

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`.

Olya, M. H., H. Badri, S. Teimoori, and K. Yang (2022). An integrated deep learning and stochastic optimization approach for resource management in team-based healthcare systems. *Expert Systems with Applications 187*, 115924.

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.

Rais, A. and A. Viana (2011). Operations research in healthcare: A survey. *International Transactions in Operational Research 18*(1), 1–31.

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*, 100031.

Reinsch, R. and L. Xiu (2025). Revamping the police: A trust-based approach to serving the us community. *Public Integrity*, 1–18.

Ren, Y., Y. Ding, and S. Zhou (2006). A data mining approach to study the significance of nonlinearity in multistation assembly processes. *IIE Transactions 38*(12), 1069–1083.

Rezaei, J., A. Arab, and M. Mehregan (2022). Equalizing bias in eliciting attribute weights in multiattribute decision-making: Experimental research. *Journal of Behavioral Decision Making 35*(2), e2262.

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.

Su, L., Y. Zhang, and J. Wei (2016). A practical test for strict exogeneity in linear panel data models with fixed effects. *Economics Letters 147*, 27–31.

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 (2023). 2020 census urban and rural classification and urban area criteria. 2020 Census.

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). Local government services. `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. and H. F. Thomas (2026). The other side of chaos COVID-19, federal spending, and local government volatility. *Journal of European Public Policy 33*(1), 373–397.

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: Data Summary

Variable (unit)	Mean	Min.	Max.	Std. dev.	Source
CP+ HS (%)	6.04	0.00	47.92	7.93	WV State Auditor
CR+GG+PS (%)	93.64	51.97	100.00	8.06	WV State Auditor
SS (%)	0.32	0.00	5.39	0.71	WV State Auditor
pop (persons)	33,053.94	5,066.00	192,248.00	32,980.76	U.S. Census Bureau
uninsured_pct (%)	9.70	3.20	23.30	4.07	ACS
med_hh_inc ($)	42,853.96	22,252.00	93,744.00	8,917.26	ACS
own30_pct (%)	14.75	3.80	31.20	3.29	ACS
pub_assist_pct (%)	2.63	0.19	11.30	1.40	ACS

Table OS.1: Descriptive statistics for the explanatory variables across all 55 counties in WV over the 2012–2022.

County	CP+HS	CR+GG+PS	SS
Boone	0.00	99.99	0.01
Cabell	0.40	99.60	0.00
Fayette	0.06	99.75	0.18
Kanawha	0.80	99.20	0.00
Lincoln	1.22	98.78	0.00
Logan	0.93	98.76	0.32
McDowell	0.00	100.00	0.00
Mercer	2.07	97.68	0.25
Mingo	0.00	100.00	0.00
Raleigh	3.24	96.34	0.42
Wayne	3.10	96.89	0.00
Wyoming	3.41	95.49	1.10

Table OS.2: Observed *expenditure shares* (%) by category for each of the 12 southern WV counties in 2022.

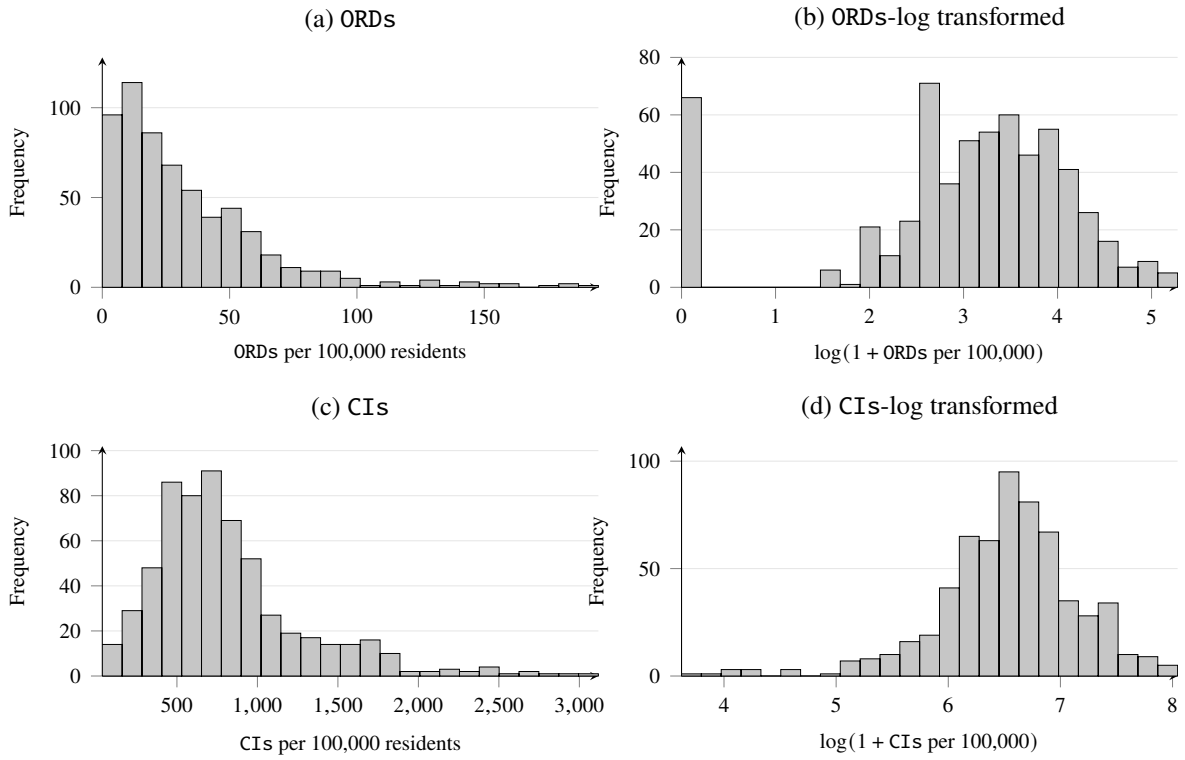


Figure OS.1: Distributions of the observed (left) and log-transformed (right) ORDs and CIs across all 55 WV counties over 2012-2022.

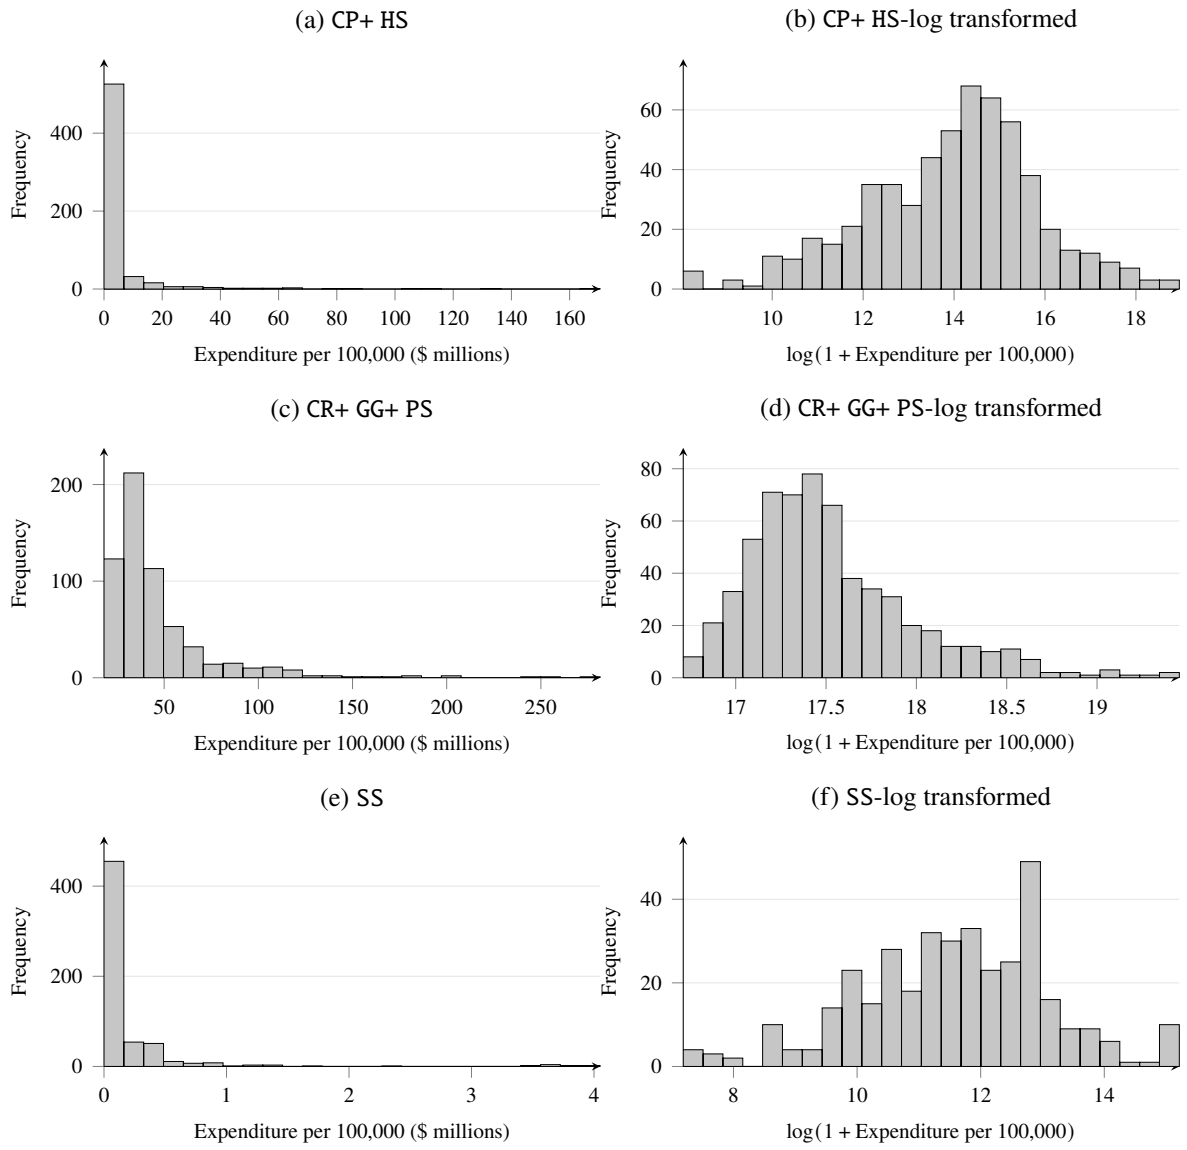
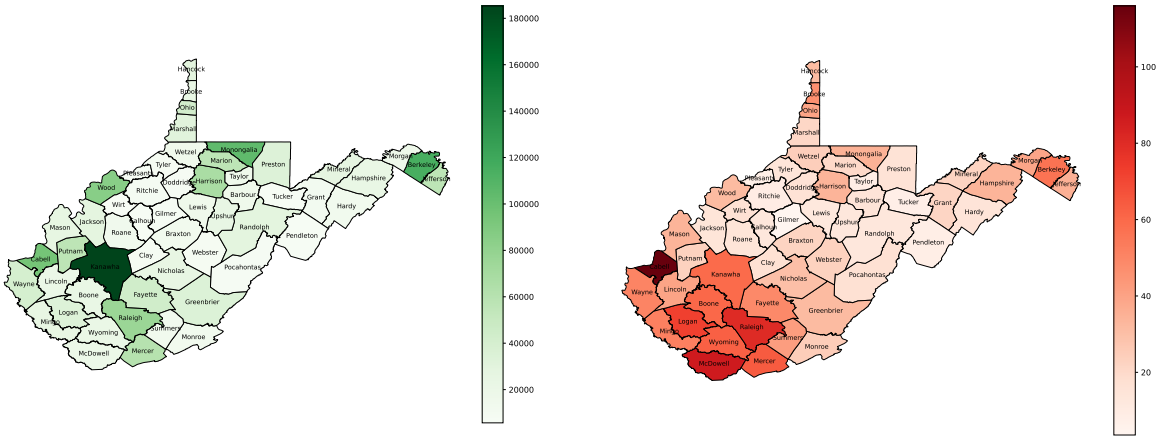


Figure OS.2: Distributions of observed expenditures (left) and their log-transformed values (right) for the three AECs across all 55 WV counties over the 2012–2022.

(a) Population

(b) ORDs per 100,000



(c) CIs per 100,000

(d) Budget per 100,000

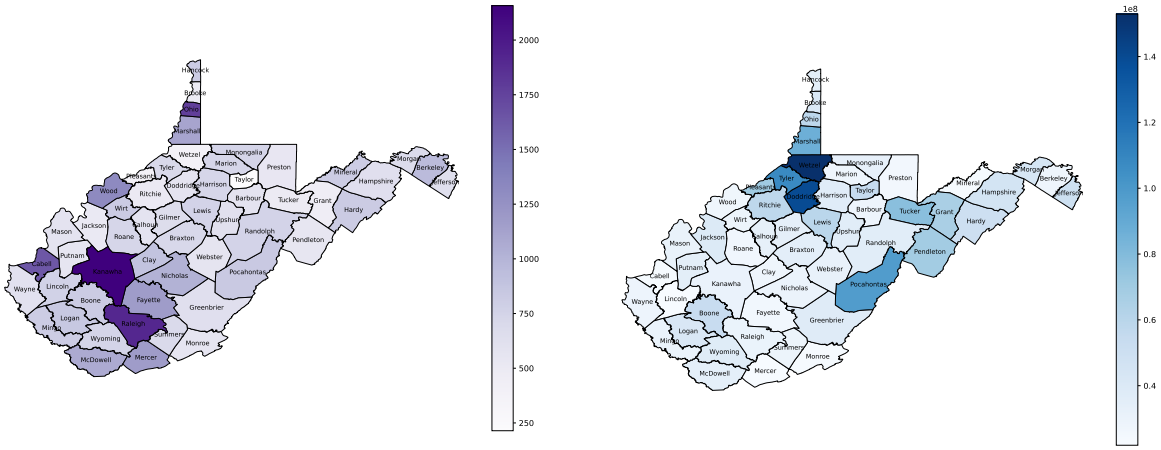


Figure OS.3: Population and average per-capita metrics across WV counties from 2012–2022. Subplots show: (a) population, (b) ORDs per 100,000, (c) CIs per 100,000, and (d) budget per 100,000.

OS.2: 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.

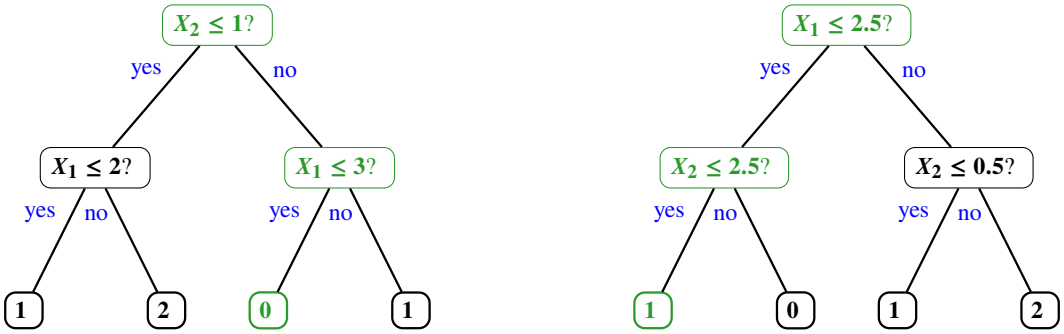


Figure OS.4: Tree ensemble predicting value $\frac{1}{2}$ (average of green leaves' values) when $(X_1, X_2) = (1.2, 2.1)$.

To give a concrete example, we represent in Figure OS.4, 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 CP+HS, X_2 as the amount money spent on SS, 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.4 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.4.

OS.3: Tuning and Evaluating the Prediction Accuracy of FE+TEs

In this section, we discuss the tuning of the TEs using the collected data and evaluate their predictive performance for ORDs and CIs in WV. Specifically, we compare the prediction accuracy of the combined FE+TEs models with that of the standalone FE and TEs models.

Table OS.3 presents the tuned hyperparameters of the TEs used to predict ORDs and CIs in the standalone TEs and combined FE+TEs models. We tune six key TE hyperparameters: (i) maximum

Model	*max_depth*		*min_samples_leaf*		*min_samples_split*		*n_estimators*		*max_features*		*max_samples*	
	ORDs	CIs	ORDs	CIs	ORDs	CIs	ORDs	CIs	ORDs	CIs	ORDs	CIs
TEs	20	10	5	5	2	2	50	50	1.0	1.0	0.8	0.8
FE+ TEs	10	20	10	5	2	2	50	150	0.7	1.0	1.0	1.0

Table OS.3: Tuned TEs hyperparameters for predicting ORDs and CIs.

tree depth (max_depth), (ii) minimum number of observations per leaf (min_samples_leaf), (iii) minimum number of observations required to split a node (min_samples_split), (iv) number of trees (n_estimators), (v) fraction of features considered at each split (max_features), and (vi) fraction of training observations used for each tree (max_samples). To assess potential spatial dependence in prediction errors, we conduct Moran's I tests on the 2023 residuals of the trained FE+TEs models across all 55 West Virginia counties. The results show no statistically significant spatial autocorrelation for either ORDs ($I = 0.0354$, $p = 0.5342$) or CIs ($I = 0.0691$, $p = 0.3042$). These findings suggest that the combined models do not exhibit significant geographic clustering of prediction errors in 2023.

	ORDs			CIs		
Model	MAE	RMSE	R^2	MAE	RMSE	R^2
FE	0.484	0.729	0.456	0.318	0.447	0.255
TEs	0.591	0.805	0.337	0.347	0.446	0.261
FE+TEs	0.458	0.690	0.513	0.319	0.439	0.282

Table OS.4: Comparison of the prediction performance of FE, TEs, and FE+TEs.

Table OS.4 compares the predictive performance of FE, TEs, and FE+TEs for ORDs and CIs in 2023 using three performance metrics: Mean Absolute Error (MAE), Root Mean Squared Error (RMSE) and Coefficient of Determination (R^2). MAE measures the average absolute difference between predicted and observed values, providing a direct measure of prediction error. RMSE is the square root of the average squared difference between predicted and observed values and, therefore, places greater weight on larger prediction errors than MAE. For both MAE and RMSE, smaller values indicate better predictive performance. The Coefficient of Determination, R^2, indicates how well a model explains the variability in the outcome variable (James et al. 2013), with higher values indicating better model fit.

As shown in Table OS.4, the FE+TEs models generally provide the strongest predictive performance for both outcomes. For ORDs, FE+TEs achieves the lowest MAE (0.458) and RMSE (0.690), as well as

the highest R^2 (0.513), outperforming both the standalone `FE` and `TEs` models across all three metrics. For `CIs`, the `FE` and `FE+TEs` models achieve similar `MAE` values of 0.318 and 0.319, respectively, while `FE+TEs` achieves the lowest `RMSE` (0.439) and the highest R^2 (0.282).

OS.4: Table of Notations for Section 5

Symbol	Definition
Sets	
I, J	Sets of outcomes and expenditure categories, respectively.
S_e	Set of socioeconomic inputs.
$\bar{J}^i, \bar{S}_e^i$	Sets of expenditure and socioeconomic inputs with significant fixed-effects coefficients for outcome i.
$\mathcal{T}_i$	Set of trees in the `FE+TEs` model trained for outcome i.
$\mathcal{L}(\tau)$	Set of leaves of tree $\tau \in \mathcal{T}_i$.
$\bar{L}_j$	Index set of expenditure intervals for input j.
$K_{l\tau}^j$	Index set of expenditure intervals for input j compatible with leaf l of tree τ.
$\mathbf{X}$	Set of feasible allocation vectors satisfying budget and `RBA` constraints.
$\bar{\mathbf{U}}_i$	Ambiguity set of admissible nondecreasing concave utility functions for outcome i.
Indices	
i, j, s	Outcome, expenditure category, and socioeconomic variable indices, respectively; $i \in I$, $j \in J$, $s \in S_e$.
τ, l	Tree and leaf indices, respectively; $\tau \in \mathcal{T}_i$, $l \in \mathcal{L}(\tau)$.
v	Expenditure interval index, $v \in \bar{L}_j$.
ι	Index of the partition of $[\theta_1^1, \theta^2]$, $\iota = 0, \ldots, w_i$.
$\tilde{c}, \tilde{t}$	County and allocation year under consideration; the outcome year is $\tilde{t} + 1$.
Parameters and fixed inputs	
b	Total available county budget per 100,000 population.
$\hat{x}_j$	Reference expenditure for expenditure category j.
κ_j^-, κ_j^+	Maximum allowable decrease and increase, respectively, relative to the reference expenditure $\hat{x}_j$ for expenditure category j.
$\hat{\mathbf{h}}$	Vector of known values of socioeconomic variables.
β_j^i, γ_s^i	Fixed effects coefficients for expenditure category j and socioeconomic variable s, respectively, for outcome i.
$\alpha_{\tilde{c}}^i, \kappa_{\tilde{t}+1}^i$	County and year fixed effects for outcome i.

Continued on next page

Table OS.5 (continued)

Symbol	Definition
$c_{l\tau}$	Leaf output for leaf l of tree τ .
v_τ	Weight of tree τ.
a_{jv}	Upper bound of expenditure interval v for category j.
σ_i	Priority weight assigned to outcome i.
θ_i^1, θ_i^2	Minimum and maximum predicted values attainable for outcome i under feasible allocations.
w_i	Number of intervals in the partition of $[\theta_i^1, \theta_i^2]$.
$p_{i\iota}$	Lower bound of the ι^{th} partition of $[\theta_i^1, \theta_i^2]$.
$\bar{u}_i^a, \bar{u}_i^b$	Lower and upper utility *reference functions* for outcome i.
$\tilde{u}_i(\omega, \lambda), \lambda$	*CRRA* utility function and its risk-aversion parameter.
Decision variables and functions	
x_j	Log-transformed expenditure allocated to category j.
$\xi_i(\mathbf{x}, \hat{\mathbf{h}})$	Predicted value of outcome i.
$\mathbf{\Phi}^i(\mathbf{x}, \mathbf{h})$	Residual prediction from `TEs` from outcome i .
z_i	Auxiliary variable representing the predicted value of outcome i.
$q_{l\tau}$	Binary variables indicating selection of leaf l on tree τ.
$\bar{x}_{jv}$	Binary variables indicating selection of expenditure interval v for expenditure category j.
$\bar{u}_i$	Non-decreasing concave utility function for outcome i.
$\bar{\pi}_i(\cdot)$	Worst-case utility for outcome i over $\bar{\mathbf{U}}_i$.
$u(\mathbf{z})$	Total weighted utility across all outcomes, $\sum_{i\in I} \sigma_i \bar{u}_i(\cdot)$.
$\hat{\mathbf{y}}_i^{1:7}$	Concatenated vector of dual variables of (9) i.

Table OS.5: Summary of key notations used for deriving $\mathcal{F}$.

OS.5: Extensive Form of $\mathcal{F}$

The budget allocation model $\mathcal{F}$ we proposed in Section 5.3 is obtained by reformulating (2). 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.

$$(\mathcal{F}) \quad \max \sum_{i\in I} \sigma_i \left[\begin{array}{l} \sum_{\iota=0}^{w_i} \left(\tilde{u}_i(p_{i\iota}, 0.41)\hat{y}^4_{i\iota} + \tilde{u}_i(p_{i\iota}, 0.68)\hat{y}^5_{i\iota} \right) \\ + \sum_{\iota=0}^{w_i-1} \left(\tilde{u}'_i(p_{i\iota}, 0.41)\hat{y}^6_{i\iota} + \tilde{u}'_i(p_{i\iota}, 0.68)\hat{y}^7_{i\iota} \right) \end{array} \right] \tag{A1}$$

$$-\hat{y}^1_{i0} - \hat{y}^2_{i0} - \hat{y}^3_{i0} + \hat{y}^4_{i0} + \hat{y}^5_{i0} \le 0 \qquad \forall i \in I \tag{A2}$$

$$\hat{y}^1_{i(\iota-1)} - \hat{y}^1_{i\iota} + \hat{y}^2_{i(\iota-1)} - \hat{y}^2_{i\iota} - \hat{y}^3_{i\iota} + \hat{y}^4_{i\iota} + \hat{y}^5_{i\iota} \le 0 \qquad \forall i \in I, \quad \iota = 1, \dots, w_i - 2 \tag{A3}$$

$$\hat{y}^1_{i(\iota-1)} - \hat{y}^1_{i\iota} + \hat{y}^2_{i(\iota-1)} - \hat{y}^3_{i\iota} + \hat{y}^4_{i\iota} + \hat{y}^5_{i\iota} \le 0 \qquad \forall i \in I, \quad \iota = w_i - 1 \tag{A4}$$

$$\hat{y}^1_{i(w_i-1)} - \hat{y}^3_{iw_i} + \hat{y}^4_{iw_i} + \hat{y}^5_{iw_i} \le 0 \qquad \forall i \in I \tag{A5}$$

$$-\hat{y}^1_{i0}(p_{i1} - p_{i0}) + \hat{y}^6_{i0} + \hat{y}^7_{i0} \le 0 \qquad \forall i \in I \tag{A6}$$

$$\begin{array}{l} -\hat{y}^1_{i\iota}(p_{i(\iota+1)} - p_{i\iota}) - \hat{y}^2_{i(\iota-1)}(p_{i\iota} - p_{i(\iota-1)}) \\ \quad + \hat{y}^6_{i\iota} + \hat{y}^7_{i\iota} \le 0 \end{array} \qquad \forall i \in I, \quad \iota = 1, \dots, w_i - 1 \tag{A7}$$

$$\sum_{\iota=0}^{w_i} \hat{y}^3_{i\iota} \le 1 \qquad \forall i \in I \tag{A8}$$

$$\begin{array}{l} \sum_{\iota=0}^{w_i} p_{i\iota}\hat{y}^3_{i\iota} \le \sum_{j\in\bar{J}^i} \beta^i_j x_j + \sum_{s\in\bar{S}^i_e} \gamma^i_s \hat{h}_s + \alpha^i_{\bar{c}} + \kappa^i_{\bar{t}+1} \\ \qquad\qquad + \sum_{\tau\in\mathcal{T}_i} \sum_{l\in\mathcal{L}(\tau)} c_{l\tau} v_\tau q_{l\tau} \end{array} \qquad \forall i \in I \tag{A9}$$

$$\sum_{j\in J} (e^{x_j} - 1) \le b \tag{A10}$$

$$\log\left(\max(1, 1 + \hat{x}_j - \kappa^-_j)\right) \le x_j \le \log(1 + \hat{x}_j + \kappa^+_j) \qquad \forall j \in J \tag{A11}$$

$$\sum_{v\in\bar{L}_j} a_{j(v-1)}\bar{x}_{jv} \le x_j \le \sum_{v\in\bar{L}_j} a_{jv}\bar{x}_{jv} \qquad \forall j \in J \tag{A12}$$

$$\sum_{v\in\bar{L}_j} \bar{x}_{jv} = 1 \qquad \forall j \in J \tag{A13}$$

$$\sum_{\substack{l\in\mathcal{L}(\tau):\\ K^j_{l\tau}\subseteq\{v_1,\dots,v_2\}}} q_{l\tau} \le \sum_{v=v_1}^{v_2} \bar{x}_{jv} \qquad \begin{array}{l} \forall i \in I, \quad \tau \in \mathcal{T}_i, \\ \forall j \in J, \quad v_1, v_2 \in \bar{L}_j,\ v_1 \le v_2 \end{array} \tag{A14}$$

$$\sum_{l\in\mathcal{L}(\tau)} q_{l\tau} = 1 \qquad \forall i \in I, \quad \tau \in \mathcal{T}_i \tag{A15}$$

$$\hat{\mathbf{y}}^1_i, \hat{\mathbf{y}}^5_i, \hat{\mathbf{y}}^7_i \le 0 \qquad \forall i \in I \tag{A16}$$

$$\hat{\mathbf{y}}^2_i, \hat{\mathbf{y}}^3_i, \hat{\mathbf{y}}^4_i, \hat{\mathbf{y}}^6_i \ge 0 \qquad \forall i \in I \tag{A17}$$

$$\mathbf{q}_{\cdot\tau} \in \{0,1\}^{|\mathcal{L}(\tau)|} \qquad \forall i \in I, \quad \tau \in \mathcal{T}_i \tag{A18}$$

$$\bar{\mathbf{x}}_j \in \{0,1\}^{|\bar{L}_j|} \qquad \forall j \in J \tag{A19}$$

$$\mathbf{x} \in \mathbb{R}^{|J|}. \tag{A20}$$

OS.6: Additional Optimization Results

δ	BPP	CPP	OPP
0.02	3.84 (14.45)	5.06 (28.58)	3.88 (14.70)
0.05	18.08 (57.93)	18.11 (54.00)	19.50 (73.16)
0.10	63.45 (230.85)	62.71 (244.82)	59.68 (229.55)

Table OS.6: Mean (maximum) computational time (seconds) for $\mathcal{F}$ by *priority policy* and δ.

County	Policy	δ	CP+HS(%)	CR+GG+PS(%)	SS (%)	ρ_c
Boone	Observed	–	0.00	99.99	0.01	1.000
	BPP	0.02	2.83	97.11	0.00	0.999
		0.05	2.83	94.11	0.00	0.969
		0.10	2.83	89.11	0.00	0.919
	OPP	0.02	2.43	97.57	0.00	1.000
		0.05	4.15	95.83	0.00	1.000
		0.10	4.24	89.11	0.00	0.933
	CPP	0.02	1.94	97.11	0.96	1.000
		0.05	5.02	94.11	0.82	0.999
		0.10	8.13	89.11	2.77	1.000
Cabell	Observed	–	0.40	99.60	0.00	1.000
	BPP	0.02	2.32	97.68	0.00	1.000
		0.05	3.31	95.12	0.00	0.984
		0.10	4.91	89.68	2.61	0.972
	OPP	0.02	0.70	98.39	0.00	0.991
		0.05	5.04	94.68	0.00	0.997
		0.10	5.04	89.68	0.00	0.947
	CPP	0.02	1.73	97.68	0.59	1.000
		0.05	3.10	95.75	0.59	0.994
		0.10	9.16	89.68	1.12	1.000

Continued on next page

Table OS.7 – continued

County	Policy	δ	CP+HS(%)	CR+GG+PS(%)	SS (%)	ρ_c
Fayette	Observed	–	0.06	99.75	0.18	1.000
	BPP	0.02	2.58	95.56	0.00	0.981
		0.05	2.58	92.56	0.00	0.951
		0.10	2.58	87.97	0.00	0.905
	OPP	0.02	4.28	95.56	0.00	0.998
		0.05	4.51	92.56	0.00	0.971
		0.10	4.51	87.97	0.00	0.925
	CPP	0.02	2.77	95.59	1.04	0.994
		0.05	6.81	92.56	0.63	1.000
		0.10	8.84	87.56	3.48	0.999
Kanawha	Observed	–	0.80	99.20	0.00	1.000
	BPP	0.02	2.58	96.82	0.00	0.994
		0.05	3.99	93.88	0.00	0.979
		0.10	3.99	88.84	0.00	0.928
	OPP	0.02	3.05	96.82	0.00	0.999
		0.05	4.11	93.88	0.00	0.980
		0.10	4.04	88.82	0.00	0.929
	CPP	0.02	1.89	96.82	0.92	0.996
		0.05	4.79	93.82	0.92	0.995
		0.10	1.70	88.82	9.48	1.000
Lincoln	Observed	–	1.22	98.78	0.00	1.000
	BPP	0.02	2.94	97.06	0.00	1.000
		0.05	5.82	94.06	0.00	0.999
		0.10	7.26	89.06	0.00	0.963
	OPP	0.02	2.94	97.06	0.00	1.000
		0.05	5.94	94.06	0.00	1.000
		0.10	7.45	89.06	0.00	0.965
	CPP	0.02	1.58	97.06	1.36	1.000
		0.05	3.42	94.06	1.71	0.992
		0.10	8.67	89.06	1.71	0.995

Continued on next page

Table OS.7 – continued

County	Policy	δ	CP+HS(%)	CR+GG+PS(%)	SS (%)	ρ_c
Logan	Observed	–	0.93	98.76	0.32	1.000
	BPP	0.02	3.05	96.95	0.00	1.000
		0.05	3.85	93.95	0.00	0.978
		0.10	3.85	88.95	0.00	0.928
	OPP	0.02	3.05	96.95	0.00	1.000
		0.05	3.31	93.95	0.00	0.973
		0.10	3.31	88.95	0.00	0.923
	CPP	0.02	2.89	96.95	0.16	1.000
		0.05	3.85	95.35	0.17	0.994
		0.10	3.85	89.98	0.17	0.940
McDowell	Observed	–	0.00	100.00	0.00	1.000
	BPP	0.02	1.79	98.21	0.00	1.000
		0.05	3.89	95.00	0.00	0.989
		0.10	3.89	90.89	0.00	0.948
	OPP	0.02	0.66	98.00	0.00	0.987
		0.05	4.01	95.00	0.00	0.990
		0.10	4.01	90.89	0.00	0.949
	CPP	0.02	1.74	98.21	0.05	1.000
		0.05	4.67	95.00	0.05	0.997
		0.10	7.94	90.20	0.91	0.991
Mercer	Observed	–	2.07	97.68	0.25	1.000
	BPP	0.02	5.58	94.16	0.00	0.997
		0.05	7.59	91.16	0.00	0.987
		0.10	7.59	86.16	0.00	0.937
	OPP	0.02	5.58	94.16	0.00	0.997
		0.05	7.59	91.16	0.00	0.987
		0.10	11.91	86.16	0.00	0.981
	CPP	0.02	5.51	94.16	0.33	1.000
		0.05	8.51	91.16	0.33	1.000
		0.10	9.73	86.16	0.33	0.962

Continued on next page

Table OS.7 – continued

County	Policy	δ	CP+HS(%)	CR+GG+PS(%)	SS (%)	ρ_c
Mingo	Observed	–	0.00	100.00	0.00	1.000
	BPP	0.02	5.53	91.67	0.00	0.972
		0.05	3.74	88.67	0.00	0.924
		0.10	3.74	83.67	0.00	0.874
	OPP	0.02	5.71	91.67	0.00	0.974
		0.05	5.71	88.67	0.00	0.944
		0.10	5.71	86.65	0.00	0.924
	CPP	0.02	6.64	91.67	0.90	0.992
		0.05	10.20	88.88	0.90	1.000
		0.10	10.20	83.67	1.29	0.952
Raleigh	Observed	–	3.24	96.34	0.42	1.000
	BPP	0.02	4.89	95.11	0.00	1.000
		0.05	7.27	91.64	0.00	0.989
		0.10	8.86	86.64	0.00	0.955
	OPP	0.02	4.89	95.11	0.00	1.000
		0.05	6.42	91.64	0.00	0.981
		0.10	9.88	86.64	0.00	0.965
	CPP	0.02	4.42	94.64	0.94	1.000
		0.05	7.27	91.64	0.94	0.998
		0.10	7.27	86.64	0.94	0.948
Wayne	Observed	–	3.10	96.89	0.00	1.000
	BPP	0.02	3.08	95.05	0.00	0.981
		0.05	3.08	92.05	0.00	0.951
		0.10	2.44	87.48	0.00	0.899
	OPP	0.02	3.54	95.94	0.00	0.995
		0.05	3.54	92.05	0.00	0.956
		0.10	3.54	88.93	0.00	0.925
	CPP	0.02	4.10	95.05	0.85	1.000
		0.05	4.20	92.06	0.88	0.971
		0.10	2.44	87.40	10.00	0.998

Continued on next page

Table OS.7 – continued

County	Policy	δ	CP+HS(%)	CR+GG+PS(%)	SS (%)	ρ_c
Wyoming	Observed	–	3.41	95.49	1.10	1.000
	BPP	0.02	2.59	97.32	0.00	0.999
		0.05	2.59	97.32	0.00	0.999
		0.10	3.65	86.38	0.00	0.900
	OPP	0.02	2.68	97.32	0.00	1.000
		0.05	2.68	97.32	0.00	1.000
		0.10	2.68	97.32	0.00	1.000
	CPP	0.02	4.61	94.38	0.28	0.993
		0.05	7.80	91.41	0.28	0.995
		0.10	12.80	86.74	0.28	0.998

Table OS.7: Observed and optimized *expenditure shares* (%) by county, *priority policy*, and δ.