—Preliminary and incomplete—
Optimal Policy at the Margin: Sufficient Statistics from Criminal Sentencing

Robert Apel* Jim Been† Hilde Wermink‡

March 26, 2026

Abstract

We characterize optimal policy when treatment effects are heterogeneous. We show that the marginal policy-relevant treatment effect (MPRTE) is a sufficient statistic for optimal policy: interventions should expand only if the marginal response improves outcomes. This result holds even when policymakers are uncertain about which individuals are affected by a reform. We apply the framework to custodial sentencing using quasi-random assignment of cases to judges and marginal treatment effect methods. The estimated marginal response is criminogenic, implying that expanding incarceration increases recidivism at the policy margin. Optimal policy therefore depends on marginal, not average, causal effects.

Keywords: Optimal Policy; Marginal Treatment Effects; Sufficient Statistics; Criminal Sentencing; Causal Inference.
JEL codes: C26, D61, K42, I38, H21.

*Rutgers School of Criminal Justice, Rutgers University-Newark.
†Department of Economics, Leiden Law School, Leiden University and Netspar.
‡Department of Criminology, Leiden Law School, Leiden University.

1 Introduction

Optimal policy depends on how individuals at the margin respond to interventions. When treatment effects are heterogeneous, average causal effects do not determine whether a policy should expand or contract. Instead, the relevant object is the response of those individuals whose treatment status changes when policy changes. This paper shows that, in a broad class of environments, the marginal policy-relevant treatment effect (MPRTE) is the sufficient statistic governing optimal policy.

Our main contribution is conceptual. We show that optimal policy can be characterized using a single estimable object: the marginal response induced by a small policy change. Specifically, policy should expand if and only if the MPRTE is favorable. This result provides a transparent link between economic theory and causal inference, reducing a potentially complex welfare problem with heterogeneous responses to a tractable sufficient-statistics representation.

We further show that this result extends to environments with policy uncertainty. Policymakers typically do not know which individuals will be affected by a reform. We formalize this uncertainty and demonstrate that the MPRTE remains sufficient for determining the direction of optimal policy changes. While higher-order features of the treatment effect distribution matter for policy magnitude, they do not affect the sign of the optimal policy adjustment. This robustness substantially broadens the applicability of the sufficient-statistics result.

Our framework builds on and connects two strands of the literature. In the rational choice model of crime (Becker, 1968), optimal policy balances deterrence and social costs at the margin. In the marginal treatment effect (MTE) framework (Heckman and Vytlacil, 2005; Florens et al., 2008; Carneiro et al., 2010), causal effects vary across individuals according to their resistance to treatment, and policy-relevant parameters are weighted averages of these heterogeneous effects. We show that these insights can be unified: the MPRTE—the derivative of policy-relevant treatment effects with respect to policy—provides the exact object needed for optimal policy design.

Our result also connects to the sufficient-statistics approach to policy evaluation (Saez, 2001; Chetty, 2009; Hendren, 2016, 2020), which shows that optimal policy can often be characterized using a small number of empirically measurable objects. In that literature, sufficient statistics typically summarize behavioral responses through elasticities or reduced-form parameters. In contrast, we derive the sufficient statistic directly from the structure of heterogeneous treatment effects. The marginal policy-relevant treatment effect (MPRTE) emerges as the object governing optimal policy because it captures the behavioral response of individuals at the policy margin. This distinction is important: rather than positing a sufficient statistic based on reduced-form responses, we show that it arises naturally as the derivative of policy-relevant treatment effects within the marginal treatment effect framework.

We implement this framework in the context of custodial sentencing. Using quasi-random assignment of criminal cases to judges, we exploit variation in judicial sentencing tendencies to identify the marginal treatment effect of incarceration across offenders. We then map these estimates into policy-relevant treatment effects and evaluate optimal policy. The estimated MTE is increasing over the empirically identified region, implying that marginal offenders experience more adverse effects of incarceration than infra-marginal offenders. The resulting MPRTE is

strictly positive, indicating that expanding custody increases recidivism at the policy margin.

The empirical application illustrates the broader contribution of the paper. While judge-IV designs have been widely used to estimate causal effects of legal decisions (e.g., Kling, 2006; Loeffler, 2013; Maestas et al., 2013; Aizer and Doyle Jr, 2015; Dobbie et al., 2018; Harding et al., 2018; Augustine et al., 2022; Bhuller et al., 2020; Wermink et al., 2024),[1] these estimates are typically not translated into explicit policy rules. We show how heterogeneous treatment effect estimates can be directly mapped into optimal policy using the MPRTE as a sufficient statistic.

More generally, the paper demonstrates how modern causal inference tools can deliver simple and robust policy guidance. By focusing on marginal responses rather than average effects, policymakers can align empirical evidence with the economic logic of optimal decision-making. The sufficient-statistics result we derive is not specific to criminal sentencing. It applies to any policy environment in which treatment assignment operates through a threshold rule and treatment effects are heterogeneous. Such settings arise in many areas of economics, including education (e.g., admission or remediation thresholds), labor markets (e.g., training program eligibility), health policy (e.g., insurance coverage rules), and public finance (e.g., benefit eligibility cutoffs). In such settings, optimal policy depends only on the marginal response induced by a policy change. This generality highlights that the MPRTE is not merely a policy parameter, but a fundamental object linking causal inference to optimal policy design.

2 Rational-choice model

This section provides a rational-choice framework that motivates why the *net* effect of custody on reoffending can be positive or negative and why it can vary across defendants. The empirical focus is on estimating this net response function and its variation with resistance, not on decomposing it into deterrence and post-sentence components.

2.1 Theory of optimal deterrence

In the economic model of crime (Becker, 1968), an individual commits an offense when the net benefit of offending exceeds expected punishment,

$$B_i - p \cdot F \geq 0.$$

with B the benefits of crime for individual i, p the probability of punishment, and F the severity of punishment. This inequality defines a latent index determining the offending decision: offenders at the margin are those for whom a small change in sanctions (either p or F) reverses the sign of the expression.[2]

The rational-choice model we develop below adopts the same threshold structure but allows the net benefit to vary across individuals through an unobserved resistance type. The latent offending index therefore corresponds to Becker's net benefit term, and custodial sanctions shift this index. The behavioral response to punishment is given by the derivative of the offending decision with respect to the sanction, which—in the heterogeneous setting—varies across the

[1] For an overview of Judge-IV papers, their context, and their results, we refer to Table 6 in the Appendix.

[2] Becker's main takeaway was to set a high F and low p for cost-effective judicial policy.

resistance distribution. The model below formalizes this representation and derives the resulting marginal response function that will serve as the policy-relevant object.

Before moving to the rational-choice model below, it is good to acknowledge what the model does and does not entail. In this paper's perspective—in line with Becker (1968)—the sole goal of sentencing is deterrence. However, criminologists have also considered....

[@Hilde/Bob: Add a paragraph on OTHER goals of sentencing.]

2.2 Reoffending opportunities

At $t = 0$, a defendant is sentenced for the index offense and receives $D \in \{0, 1\}$, where $D = 1$ denotes custody and $D = 0$ a noncustodial sanction. After sentencing, the individual faces opportunities to reoffend at dates $t + s$, $s \in \{1, \ldots, n\}$. Let $a_{is} \in \{0, 1\}$ indicate offending at horizon s. Utility from offending is

$$V_{is}(1; D) = b_{is} - \pi_{is}(D)\, s_{is} - \kappa_{is}(D) + \varepsilon_{is}, \tag{1}$$

and utility from not offending is normalized to $V_{is}(0; D) = 0$. Here b_{is} is the private benefit of offending; $\pi_{is}(D)$ is perceived detection/conviction risk; s_{is} is expected sanction severity upon detection; and $\kappa_{is}(D)$ captures post-sentence constraints that affect the net payoff to reoffending but are not modeled as belief updates. Shocks ε_{is} are continuous.

Define recidivism as any reconviction during the horizon:

$$Y_i(D) = \mathbb{1}\left\{\exists s \in \{1, \ldots, n\} : V_{is}(1; D) \geq 0\right\}. \tag{2}$$

The causal effect of custody is $\tau_i \equiv Y_i(1) - Y_i(0)$.

2.3 Threshold proximity and local responsiveness

Define the latent index $\eta_{is}(D) \equiv b_{is} - \pi_{is}(D)\, s_{is} - \kappa_{is}(D)$ so that $a_{is}(D) = \mathbb{1}\{\eta_{is}(D) + \varepsilon_{is} \geq 0\}$. Let F and f denote the CDF and PDF of ε_{is}.

Lemma 1 (Local approximation)**.** *Suppose ε_{is} has density f and custody shifts the index by $\Delta_{is} \equiv \eta_{is}(1) - \eta_{is}(0)$. Then*

$$\Pr(a_{is} = 1 \mid D = 1) - \Pr(a_{is} = 1 \mid D = 0) = F(\eta_{is}(1)) - F(\eta_{is}(0)) \approx f(\eta_{is}(0)) \cdot \Delta_{is},$$

for small shifts. Net behavioral effects are therefore larger in magnitude for individuals closer to the offending threshold (higher $f(\eta_{is}(0))$).

2.4 A bridge to MTE

Define per-period probability shifts $\delta_{is} \equiv F(\eta_{is}(1)) - F(\eta_{is}(0)) \approx f(\eta_{is}(0))\Delta_{is}$. Because $Y_i(D)$ is an "any-offense" event across horizons, τ_i aggregates across s. A transparent approximation

is obtained by linearizing the event probability:

$$\tau_i \;\approx\; \sum_{s=1}^{n} w_{is}\,\delta_{is} \;\approx\; \sum_{s=1}^{n} w_{is}\,f(\eta_{is}(0))\,\Delta_{is}, \tag{3}$$

where weights $w_{is} \in (0,1]$ reflect that changes in earlier offending probabilities affect the probability of any offense over the horizon. Equation (3) will map directly into the MTE as a conditional expectation of the net structural response given resistance rank.

3 MTE and estimands

3.1 Selection and resistance rank

Custody follows a latent-index selection rule:

$$D_i = \mathbb{1}\{P(Z_i, X_i) \geq R_i\}, \tag{4}$$

where Z is judge leniency, X are covariates, $P(Z,X) = \Pr(D = 1 \mid Z, X)$ is the propensity score, and R_i is unobserved resistance to custody. Let F_R denote the CDF of R. Define

$$U_i \equiv F_R(R_i). \tag{5}$$

By the probability integral transform, if R is continuous then $U \sim \text{Unif}(0,1)$. Lowercase u denotes a realization of U.

3.2 MTE as a structural response function

The marginal treatment effect is

$$\text{MTE}(u) \equiv \mathbb{E}[Y_i(1) - Y_i(0) \mid U_i = u] = \mathbb{E}[\tau_i \mid U_i = u]. \tag{6}$$

Combining (6) with the approximation in (3) yields the structural mapping:

$$\text{MTE}(u) \;\approx\; \mathbb{E}\left[\sum_{s=1}^{n} w_{is}\,f(\eta_{is}(0))\,\Delta_{is} \;\middle|\; U_i = u\right]. \tag{7}$$

Thus, MTE(u) is interpretable as a structural (rationalized) net response function: it is the conditional average of the rational-choice-implied net effect of custody on reoffending.

3.3 Integrated estimands and their weights

An *estimand* is a target causal parameter defined by counterfactual outcomes $(Y(1), Y(0))$ and a population of interest. In the MTE framework, canonical estimands are integrals of MTE(u) over $u \in [0,1]$ with estimand-specific weights (Heckman and Vytlacil, 2005; Carneiro et al., 2010).

We focus on ATE, ATT, ATUT, and LATE. Using weights is not merely notational: it makes explicit which regions of the resistance distribution each parameter emphasizes. (Equivalently,

one can write some estimands as integrals over sub-intervals only under special cases such as a constant propensity score; with covariate-dependent propensities, weights are the natural representation.)

Formally,

$$\text{ATE} = \int_0^1 \text{MTE}(u)\, du, \tag{8}$$

$$\text{ATT} = \int_0^1 \text{MTE}(u)\, w_{\text{ATT}}(u)\, du, \tag{9}$$

$$\text{ATUT} = \int_0^1 \text{MTE}(u)\, w_{\text{ATUT}}(u)\, du, \tag{10}$$

$$\text{LATE} = \int_0^1 \text{MTE}(u)\, w_{\text{IV}}(u)\, du. \tag{11}$$

What are the weights? Let $P \equiv P(Z, X)$ denote the propensity score and note that $\Pr(D = 1) = \mathbb{E}[P]$ under the selection model. A convenient representation (using $f_U(u) = 1$ for $U \sim \text{Unif}(0, 1)$) is:

$$w_{\text{ATT}}(u) = f_{U|D=1}(u) = \frac{\Pr(D = 1 \mid U = u)\, f_U(u)}{\Pr(D = 1)} = \frac{\Pr(P \geq u)}{\mathbb{E}[P]}, \tag{12}$$

$$w_{\text{ATUT}}(u) = f_{U|D=0}(u) = \frac{\Pr(D = 0 \mid U = u)\, f_U(u)}{\Pr(D = 0)} = \frac{\Pr(P < u)}{\mathbb{E}[1 - P]}. \tag{13}$$

These weights reflect selection: ATT upweights low-u (low resistance) types who are frequently treated; ATUT upweights high-u types who are frequently untreated.

For LATE, the weight depends on how the instrument shifts treatment probabilities. In the continuous-instrument representation used in the MTE literature,

$$w_{\text{IV}}(u) \;\propto\; \mathbb{E}\left[\frac{\partial P(Z, X)}{\partial Z} \;\middle|\; U = u\right], \qquad \int_0^1 w_{\text{IV}}(u)\, du = 1, \tag{14}$$

so that LATE concentrates weight on resistance ranks where judge assignment generates the greatest change in custody probabilities.

Weights versus threshold proximity. The weights $w(\cdot)$ specify *which resistance ranks are averaged*; threshold proximity in Section 2 enters through $f(\eta_{is}(0))$ inside MTE(u) in (7). The two concepts are linked because if baseline proximity and custody-induced index shifts covary with resistance rank, then MTE(u) varies with u and different estimands differ because they average the same structural response function over different regions.

3.4 Slope cases and what they imply

The slope of MTE(u) is an empirical object. Combining (7) with the weight functions clarifies implications:

- If MTE(u) is **flat**, effects are homogeneous in resistance rank and ATE, ATT, ATUT, and LATE coincide up to sampling variation.

- If MTE(u) is **downward sloping**, marginal defendants (higher u) have smaller net effects than low-resistance defendants; LATE may be smaller than ATT.
- If MTE(u) is **upward sloping**, marginal defendants have larger net effects than low-resistance defendants; LATE may exceed ATT and can differ sharply from ATE.

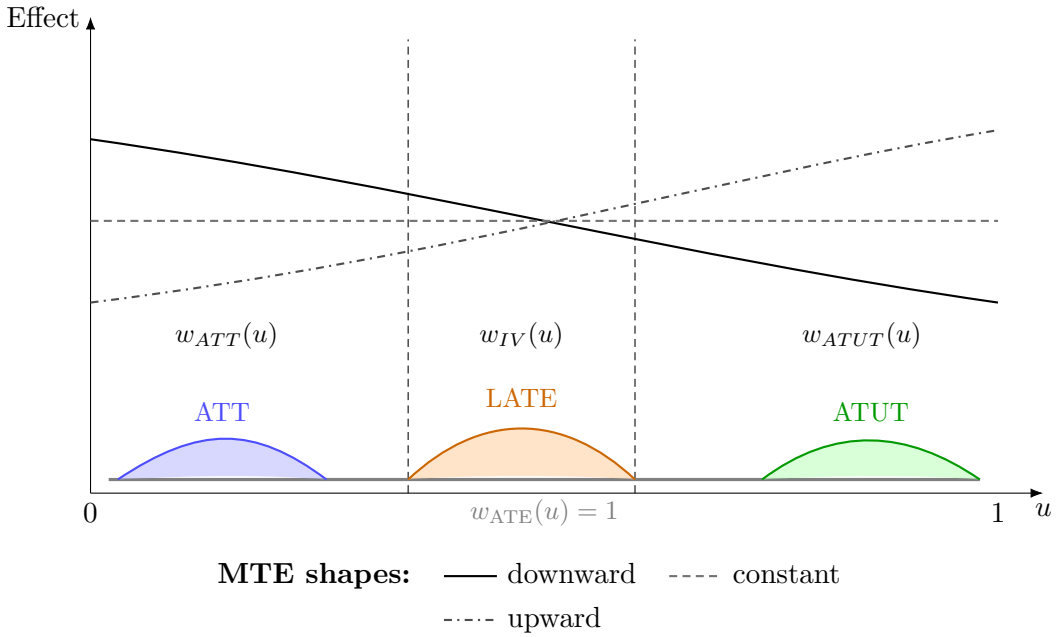


Figure 1: Schematic: MTE(u) may be downward sloping, constant, or upward sloping. ATE, ATT, ATUT, and LATE are weighted averages of the same MTE(u); the shaded weights indicate which resistance ranks are emphasized by each estimand.

3.5 PRTE and MPRTE (discrete vs marginal reforms)

Definitions and acronyms. **PRTE** is the *policy-relevant treatment effect*: the average causal effect of a *discrete* reform. **MPRTE** is the *marginal policy-relevant treatment effect*: the causal effect of an *infinitesimal* (marginal) reform. We parameterize reforms by θ and write the counterfactual propensity score as $P_\theta(Z, X)$, with baseline $P_0(Z, X)$.

Examples: (i) a *discrete* reform that reduces custody probabilities by five percentage points for a set of cases (a guideline adjustment); (ii) a *marginal* reform that slightly tightens custody eligibility (an infinitesimal change to a threshold).

Under $D = \mathbb{1}\{U \leq P(Z, X)\}$, expected outcomes under policy θ satisfy

$$\mathbb{E}[Y_\theta] = \mathbb{E}[Y(0)] + \mathbb{E}\left[\int_0^{P_\theta(Z,X)} \text{MTE}(u)\, du\right]. \tag{15}$$

Discrete reform (PRTE). The PRTE for a discrete change from P_0 to P_θ is

$$\text{PRTE}(\theta) = \frac{\mathbb{E}\left[\int_{P_0(Z,X)}^{P_\theta(Z,X)} \text{MTE}(u)\, du\right]}{\mathbb{E}[P_\theta(Z, X) - P_0(Z, X)]}. \tag{16}$$

Marginal reform (MPRTE). The MPRTE is the derivative ratio evaluated at $\theta = 0$:

$$\text{MPRTE} \equiv \left.\frac{d\mathbb{E}[Y_\theta]/d\theta}{d\mathbb{E}[D_\theta]/d\theta}\right|_{\theta=0}. \tag{17}$$

Differentiating (15) w.r.t. θ gives

$$\frac{d\mathbb{E}[Y_\theta]}{d\theta} = \mathbb{E}\left[\text{MTE}(P_\theta(Z,X)) \cdot \frac{\partial P_\theta(Z,X)}{\partial \theta}\right], \qquad \frac{d\mathbb{E}[D_\theta]}{d\theta} = \mathbb{E}\left[\frac{\partial P_\theta(Z,X)}{\partial \theta}\right],$$

and therefore

$$\text{MPRTE} = \frac{\mathbb{E}\left[\text{MTE}(P_0(Z,X)) \cdot \left.\frac{\partial P_\theta}{\partial \theta}\right|_{\theta=0}\right]}{\mathbb{E}\left[\left.\frac{\partial P_\theta}{\partial \theta}\right|_{\theta=0}\right]}. \tag{18}$$

How PRTE/MPRTE relate to ATE/ATT/ATUT/LATE. All of these objects are weighted averages of MTE(u). ATE uses uniform weights (8). ATT and ATUT use the selection weights (12)–(13). LATE uses instrument-induced weights (14). PRTE and MPRTE use *reform-induced* weights, determined by how the reform shifts $P_\theta(Z,X)$ relative to $P_0(Z,X)$ in (16)–(18). Thus, the MTE provides a single structural response function from which all estimands and reform effects follow by weighting and integration.

3.6 Hypotheses

We organize hypotheses from most basic to most transportability-focused, and we place the marginal reform hypothesis last. The MPRTE is a policy object built from the preceding knowledge about ATE and LATE.

- **H1 (Sentencing-margin sign):** LATE < 0 versus LATE > 0.
- **H2 (Marginal vs treated):** LATE $-$ ATT > 0 (consistent with upward-sloping MTE(u) over the IV margin).
- **H3 (Marginal vs untreated):** LATE $-$ ATUT < 0 (or > 0 depending on where the IV margin lies relative to the untreated mass).
- **H4 (External validity gap):** LATE $\neq$ ATE.
- **H5 (Reform derivative):** MPRTE > 0 for a marginal increase in custody.

Based on **H5**, we theoretically and empirically show that the MPRTE is a sufficient-statistic for optimal custodial policy.

4 Policy simulations: discrete and marginal reforms

H5 can be empirically tested using policy simulations. This section makes explicit that such policy exercises correspond to estimating PRTE and MPRTE for particular reform mappings $\theta \mapsto P_\theta$ (marginal reforms) and $\Delta \mapsto P^\Delta$ (discrete reforms).

4.1 Reform mappings

Discrete reforms (Δ). A discrete reform indexed by Δ changes custody probabilities by a non-negligible amount:

$$P^{\Delta}(Z, X) = \min\{1, \max\{0,\ P_0(Z, X) + \Delta\}\}. \tag{19}$$

For example, $\Delta = -0.05$ represents a five-percentage-point downward shift in custody propensity for the affected population.

Marginal reforms (θ). A marginal reform indexed by θ applies a smooth perturbation:

$$P_\theta(Z, X) = P_0(Z, X) + \theta \cdot g(Z, X), \tag{20}$$

where $g(Z, X)$ encodes which cases are targeted (e.g., $g = 1$ for all cases, or $g = \mathbb{1}\{\text{property offense}\}$). The MPRTE is evaluated at $\theta = 0$.

4.2 Estimating reform effects and testing Δ and θ

Discrete reform effect. For a given Δ, the implied change in mean outcomes is

$$\mathbb{E}[Y^{\Delta}] - \mathbb{E}[Y^{0}] = \mathbb{E}\left[\int_{P_0(Z,X)}^{P^{\Delta}(Z,X)} \mathrm{MTE}(u)\, du\right],$$

and the PRTE is

$$\mathrm{PRTE}(\Delta) = \frac{\mathbb{E}\left[\int_{P_0(Z,X)}^{P^{\Delta}(Z,X)} \mathrm{MTE}(u)\, du\right]}{\mathbb{E}[P^{\Delta}(Z, X) - P_0(Z, X)]}. \tag{21}$$

Testing a discrete reform corresponds to testing $H_0 : \mathrm{PRTE}(\Delta) = 0$ for relevant Δ values (e.g., $\Delta \in \{-0.10, -0.05, +0.05\}$). In practice, inference is obtained by repeating the entire procedure (propensity score, MTE estimation, and integration) within the judge-cluster bootstrap.

Marginal reform effect. Under the perturbation in (20), the MPRTE equals

$$\mathrm{MPRTE}(g) = \frac{\mathbb{E}[\mathrm{MTE}(P_0(Z, X)) \cdot g(Z, X)]}{\mathbb{E}[g(Z, X)]},$$

when g does not depend on θ and the derivative $\partial P_\theta / \partial \theta = g$ is constant in θ. This expression clarifies that marginal reforms are weighted averages of the MTE evaluated at baseline propensities, with weights determined by the targeting rule g. Testing a marginal reform corresponds to testing $H_0 : \mathrm{MPRTE}(g) = 0$ for a given $g(\cdot)$. Again, inference follows from the judge-cluster bootstrap, recomputing $\mathrm{MPRTE}(g)$ in each replicate.

Parameter assumptions and support. All reform calculations are restricted to the empirically identified domain: when $P^{\Delta}(Z, X)$ moves mass outside the observed support of P_0, we truncate the integration region to the overlap support and report that the reform effect is

identified only for the supported range. This discipline avoids extrapolating MTE(u) outside the region where it is estimated.

5 Data

5.1 Description of the data

[@Hilde adds results.]

5.2 Institutional setting

[@Bob/Hilde add description and cross-country comparison.]

5.3 Quasi-random judge assignment

The analysis uses above mentioned administrative records covering all criminal cases adjudicated in the above mentioned institutional setting of the Netherlands in 2012. After a case is referred to criminal court, it is assigned to a judge according to administrative scheduling procedures rather than case characteristics. Conditional on observable case attributes, this assignment generates quasi-random variation in judicial sentencing tendencies.

Judges differ systematically in their propensity to impose custodial sanctions. We exploit this variation as an instrumental variable. The identifying assumption is that, conditional on case characteristics and court-time controls, judge assignment affects subsequent offending only through its effect on the probability of receiving custody.

The instrument is constructed as leave-one-out judge stringency, defined as the fraction of custodial sentences imposed by the assigned judge in other cases within the same period. Because cases are administratively allocated across available judges, this measure captures exogenous variation in the probability of incarceration rather than endogenous sorting of offenders across judges.

5.4 Linking judge assignment to latent resistance

The judge assignment mechanism provides exogenous variation in the custody threshold of the latent index model. More stringent judges apply a higher effective custody threshold, shifting the propensity score $P(Z_i, X_i)$ upward. Offenders whose latent resistance U_i lies near this threshold change treatment status when assigned to a different judge. These individuals correspond to the marginal offenders identified by the instrument.

Formally, assignment to a stricter judge increases $P(Z_i, X_i)$ while holding (Y_{1i}, Y_{0i}, U_i) fixed, implying that observed treatment variation arises from movement along the resistance distribution rather than changes in offender characteristics. The judge instrument therefore identifies the causal effect of custody for individuals at the sentencing margin, providing the empirical analogue of the marginal decision rule in the theoretical model. The MTE curve can thus be interpreted as the structural behavioral response function governing reoffending decisions across the resistance distribution.

5.5 Summary statistics

[@Hilde, insert table with summary stats here]

6 Empirical strategy

6.1 Design

We estimate custody effects using quasi-random assignment to judges with different custody propensities. The instrument is judge leniency constructed from sentencing histories. The treatment is custody for the index offense; outcomes are post-sentence reconvictions within the follow-up horizon. The baseline specification includes offense-category controls, first-offense indicator, gender, a quadratic in age, pretrial detention, nationality and region indicators, and court and year fixed effects.

Assumption 1 (Conditional random assignment)**.** *Conditional on* $(X, \textit{court FE}, \textit{year FE})$*, assignment to judges is as-good-as random.*

Assumption 2 (Exclusion)**.** *Judge leniency affects recidivism only through its effect on custody.*

Assumption 3 (Monotonicity)**.** *Assignment to a stricter judge weakly increases the probability of custody for all defendants.*

6.2 Empirical specification

We estimate heterogeneous sentencing effects using the marginal treatment effect framework. Let Y_i denote an indicator of reoffending following disposition of the index case, and D_i an indicator equal to one if the offender receives a custodial sentence rather than a non-custodial sanction. Outcomes follow the potential outcome model

$$Y_i = Y_{0i} + D_i(Y_{1i} - Y_{0i}),$$

where Y_{1i} and Y_{0i} denote potential recidivism outcomes under custody and non-custody, respectively.

Sentencing is determined by a latent index rule

$$D_i = 1\{P(Z_i, X_i) \geq U_i\},$$

where $P(Z_i, X_i)$ is the propensity score, Z_i is judge stringency, X_i are observed case characteristics, and U_i is an unobserved resistance to custody normalized to the unit interval.

The identification assumption is conditional independence,

$$(Y_{1i}, Y_{0i}, U_i) \perp Z_i \mid X_i.$$

The marginal treatment effect is defined as

$$\text{MTE}(u) = \mathbb{E}[Y_{1i} - Y_{0i} \mid U_i = u, X_i = x].$$

Following the control-function representation, we estimate

$$Y_i = \alpha + \beta D_i + g(\hat{P}_i) + D_i \cdot h(\hat{P}_i) + X_i'\gamma + \varepsilon_i,$$

where $g(\cdot)$ and $h(\cdot)$ are flexible polynomials in the estimated propensity score $\hat{P}_i$. The derivative

$$\frac{\partial \mathbb{E}[Y_i | P_i = p]}{\partial p}$$

recovers the MTE at $u = p$. Standard treatment parameters such as ATE, ATT, and LATE are obtained as weighted averages of the MTE across the relevant regions of the resistance distribution.

Judge stringency instrument. The instrumental variable Z_i measures the sentencing tendency of the assigned judge using a leave-one-out mean. Let $j(i)$ denote the judge assigned to case i, and let $\mathcal{I}_j$ be the set of all cases handled by judge j within the same court and time period. Define

$$Z_i = \frac{1}{|\mathcal{I}_{j(i)}| - 1} \sum_{\ell \in \mathcal{I}_{j(i)},\, \ell \neq i} D_\ell,$$

where D_ℓ indicates whether custody was imposed in case ℓ. Thus Z_i equals the fraction of custodial sentences imposed by judge $j(i)$ in other cases, excluding case i itself.

The leave-one-out construction ensures that the instrument varies only with the judge's overall sentencing tendency rather than mechanically reflecting the outcome of the focal case. Under quasi-random assignment of cases to judges within court-time cells, variation in Z_i shifts the probability of custody $P(Z_i, X_i)$ while remaining orthogonal to unobserved determinants of reoffending. Consequently, Z_i provides exogenous variation in the sentencing threshold of the latent index selection equation.

We estimate the MTE using the routine developed by Andresen (2018).

6.3 Inference

We implement a cluster bootstrap at the judge level with 499 replications. Inference for MTE-based objects is nonstandard because integrated estimands and reform parameters are nonlinear functionals of estimated propensity scores and the estimated MTE curve. Resampling-based inference is therefore common in semiparametric treatment-effect settings and in designs where the instrument varies at a higher level of aggregation (Andrews and Buchinsky, 2000; Cameron and Miller, 2015). The choice of 499 replications reflects diminishing returns beyond several hundred draws and follows formal guidance for selecting the number of bootstrap replications (Andrews and Buchinsky, 2000).

7 Results

7.1 Empirical tests of the assumptions

[This section should substantively refer to the text in Sections 5.2 and 5.3 and Wermink et al. (2024).]

7.1.1 Assumption 1: Conditional random assignment

[Include Balance Check results here.]

7.1.2 Assumption 2: Exclusion

[Include Placebo Regressions for predetermined outcomes here.]

7.1.3 Assumption 3: Monotonicity

[Include subgroup-specific first-stage sign checks here.]

7.1.4 Frandsen-Lefgren-Leslie Test

[Include results for the test proposed in Frandsen et al. (2025) here.]

7.2 Empirical common support

The common-support figure (Figure 2) shows that overlap between treated and untreated observations is concentrated in the lower and middle part of the propensity-score distribution. Untreated observations are heavily concentrated near zero, whereas treated observations are more evenly distributed and extend further into higher propensity scores. This implies that the MTE is most credibly identified on the region where both groups are observed with non-negligible density. Outside that region, the estimated MTE curve relies increasingly on functional-form extrapolation rather than direct support from the data. The figure therefore justifies restricting policy interpretation to the empirically supported interval of the resistance distribution. This is further investigated below.

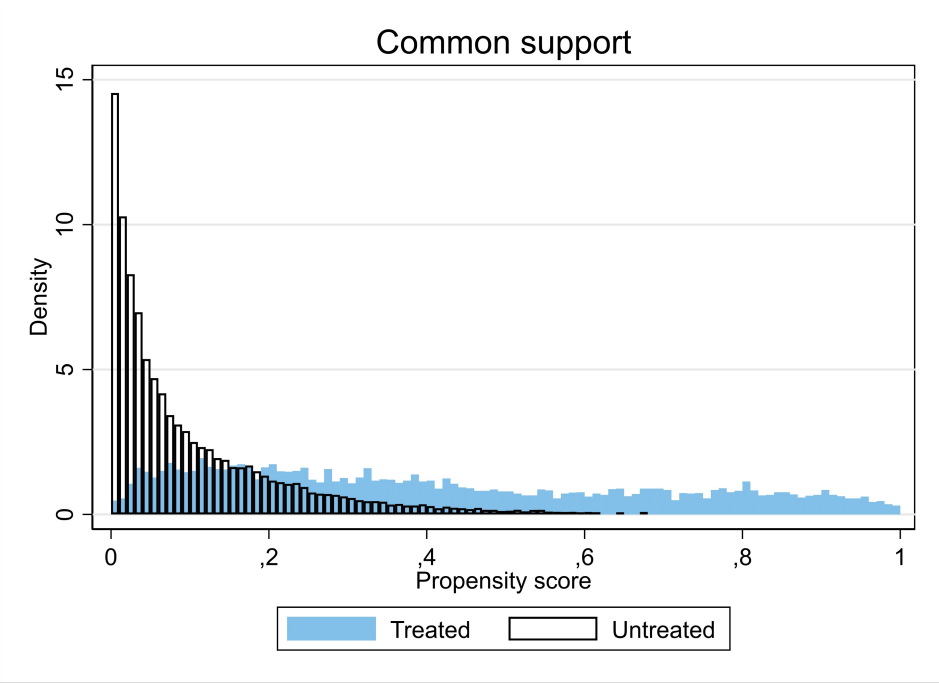


Figure 2: Common support of propensity scores for custodial versus noncustodial sentencing.

7.3 Empirical MTE-curve

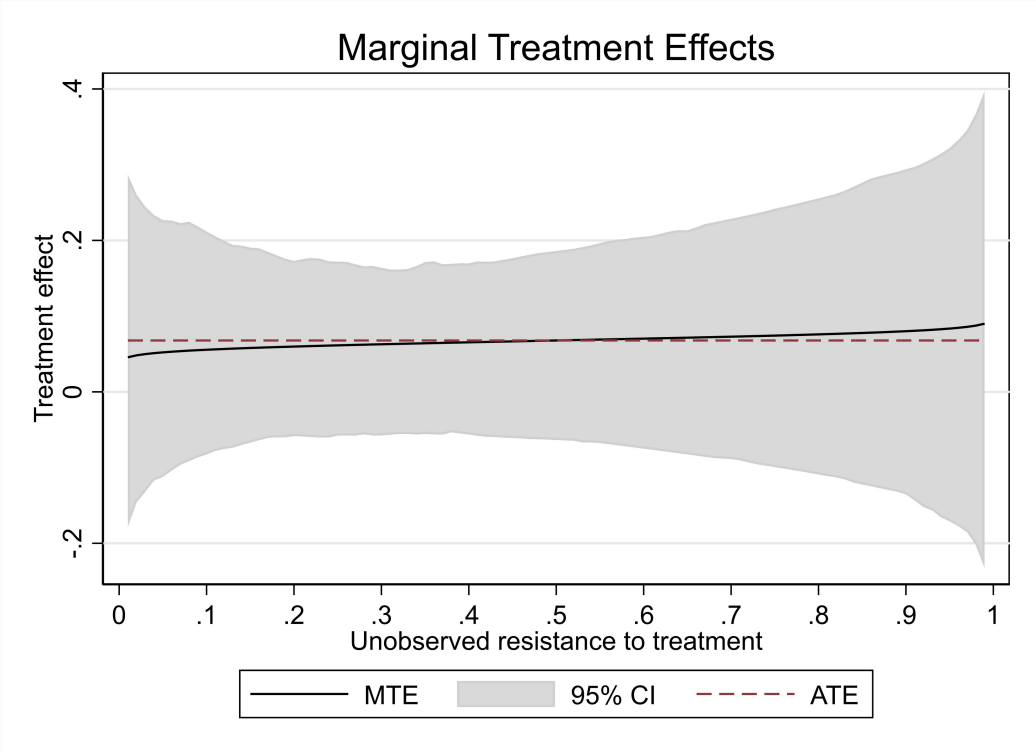


Figure 3: Estimated marginal treatment effect MTE(u) over the region of common support.

Figure 3 indicates that MTE(u) is upward sloping over the region of overlap: net custody effects are larger for higher resistance ranks. In the structural mapping (7), this implies that

$$\mathbb{E}\left[\sum_{s=1}^{n} w_{is} f(\eta_{is}(0))\Delta_{is} \mid U = u\right]$$

increases with u on the identified domain. This pattern is consistent with custody generating larger positive net index shifts and/or higher baseline proximity among higher-resistance types, and it rationalizes why LATE can exceed ATT in our estimates.

7.4 Estimation results

Table 1 reports the full set of estimates from the marginal treatment effect model. For sake of brevity, we highlight four aspects: offenders' characteristics, the strength of the judge-IV instrument, the validity of the MTE framework, and—in the next section—the implications for policy-relevant heterogeneity.

Offenders' characteristics Panel A does not identify the marginal offenders directly, but it does indicate which observable characteristics shift defendants toward or away from the custody threshold. Marginal offenders are those whose observable and unobservable characteristics place them near the treatment cutoff in the latent-index model, so that assignment to a stricter versus more lenient judge changes their incarceration status. The significant coefficients in Panel A therefore identify the observables that move defendants closer to or further from that margin.

In particular, categories associated with sizable and precisely estimated treatment-equation coefficients—in particular prior sanctions—contribute more strongly to sorting across the custody threshold. However, because marginality is defined jointly by observables and unobserved resistance, a separate complier- or local-margin characterization exercise is required to recover the observable profile of the offenders located at the policy margin. We will do so in the remainder of the paper.

Interestingly, our results suggest that there is little selection on unobservables (k (Mills)), i.e. sorting into custody is mostly driven by observables and the judge-leniency.

Panel B describes the outcome process of recidivism under incarceration and, together with the untreated equation, determines how treatment effects vary across the unobserved resistance distribution. The coefficients in Panel B describe how recidivism varies across observable characteristics among offenders who are incarcerated. A statistically significant coefficient indicates that, conditional on incarceration, offenders with that characteristic have systematically higher or lower recidivism than the reference group. Importantly, these coefficients do not represent treatment effects. They reflect differences in outcome levels under incarceration rather than the causal impact of incarceration itself. Treatment effects are obtained from the difference between outcomes under incarceration and non-incarceration and therefore depend jointly on the treated and untreated outcome equations. Panel B should thus be interpreted as characterizing the outcome process under custody, which, together with the untreated outcome equation, determines the shape of the marginal treatment effect.

The estimates further indicate the interesting fact that, among incarcerated offenders, individuals with a prior community service order (CSO) exhibit lower recidivism rates. CSO offenders are likely to be lower-risk individuals with more favorable underlying characteristics, which translates into lower recidivism even when they are incarcerated. Importantly, this does not imply that incarceration reduces recidivism for this group. The relevant object for policy is the counterfactual comparison between incarceration and non-incarceration. Our MTE estimates show that, at the policy margin, incarceration increases recidivism, implying that even for relatively low-risk offenders such as those with prior CSOs, custodial sentences may be criminogenic relative to non-custodial alternatives.

Instrument strength. Judge leniency enters the treatment equation as the instrument determining the propensity score. The results suggest....

Model diagnostics. The tests of observable and essential heterogeneity strongly reject the null of homogeneous treatment effects. This provides direct empirical support for the use of the MTE framework. In particular, essential heterogeneity implies that selection into incarceration is based on unobserved characteristics that are correlated with potential outcomes. As a result, average treatment effects do not provide a sufficient basis for policy evaluation, and a marginal analysis is required.

Table 1: Estimation results Marginal Treatment Effects

Panel A. Treatment equation

Variable	**Coef.**	**Boot. SE**
Mishandeling	-0.025^{**}	(0.013)
Diefstal 1	0.110^{**}	(0.052)
Seksueel	0.005	(0.031)
Valsheidsen	-0.020	(0.021)
Eenvoudig	0.100^{***}	(0.015)
Diefstal 2	0.015	(0.022)
Overig geweld	0.046^{**}	(0.020)
Overig vermogen	0.004	(0.015)
Misdrijven 1	-0.015	(0.017)
Misdrijven 2	0.000	(0.020)
Vernielingen	0.007	(0.015)
Verkeer	-0.069^{***}	(0.014)
Drugs	-0.037^{**}	(0.015)
Wapens en munitie	0.012	(0.026)
Overig	0.005	(0.016)
VGALG	0.000	(0.001)
VGGEV	0.020^{***}	(0.002)
VGTAAK	0.035^{***}	(0.003)
VGBOETE	0.015^{***}	(0.002)
MAXDREIG	-0.008^{***}	(0.003)
NFEIT	0.022^{***}	(0.004)
Female	-0.014^{**}	(0.007)
Age	-0.475^{***}	(0.120)
Allochtoon	0.003	(0.007)
Age^2	0.130^{***}	(0.038)
Age^3	-0.011^{***}	(0.004)
Amsterdam	0.009	(0.012)
Noord-Holland	-0.010	(0.012)
Midden-Nederland	0.013	(0.012)
Noord-Nederland	-0.029^{**}	(0.011)
Den Haag	-0.009	(0.012)
Rotterdam	0.001	(0.013)
Limburg	-0.013	(0.013)
Oost-Brabant	0.006	(0.014)
Zeeland-West	-0.004	(0.012)
Oost-Nederland (1)	-0.009	(0.012)
Oost-Nederland (2)	-0.003	(0.015)
Constant	0.711^{***}	(0.126)

Panel B. Outcome equation under treatment

Variable	**Coef.**	**Boot. SE**
Mishandeling	0.092	(0.098)
Diefstal 1	0.107	(0.145)
Seksueel	0.158	(0.183)
Valsheidsen	-0.174^{*}	(0.096)
Eenvoudig	0.110	(0.087)
Diefstal 2	0.206^{**}	(0.105)
Overig geweld	0.025	(0.093)
Overig vermogen	0.044	(0.108)
Misdrijven 1	-0.013	(0.167)
Misdrijven 2	-0.020	(0.129)
Vernielingen	0.202^{*}	(0.122)
Verkeer	-0.308^{***}	(0.104)
Drugs	0.123	(0.103)
Wapens en munitie	-0.269	(0.176)
Overig	-0.300^{**}	(0.145)
VGALG	-0.005^{***}	(0.002)
VGGEV	-0.002	(0.005)
VGTAAK	-0.018^{**}	(0.008)
VGBOETE	-0.004	(0.006)
MAXDREIG	-0.015^{*}	(0.008)
NFEIT	-0.013	(0.010)
Female	-0.027	(0.029)
Age	-0.475	(0.577)
Allochtoon	0.009	(0.031)
Age^2	0.179	(0.173)
Age^3	-0.021	(0.017)
Amsterdam	0.487	(0.331)
Noord-Holland	0.441	(0.328)
Midden-Nederland	0.462	(0.334)
Noord-Nederland	0.598^{*}	(0.333)
Den Haag	0.624^{*}	(0.329)
Rotterdam	0.541^{*}	(0.328)
Limburg	0.530	(0.330)
Oost-Brabant	0.489	(0.333)
Zeeland-West	0.405	(0.334)
Oost-Nederland (1)	0.571^{*}	(0.334)
Oost-Nederland (2)	0.533	(0.329)
Constant	0.150	(0.630)

Panel C. Selection on unobservables	**Coef.**	**Boot. SE**
k (Mills)	0.010	(0.056)
Panel D. Diagnostics		
Observable heterogeneity (p-value)	$< .001$	
Essential heterogeneity (p-value)	.026	

Note. Bootstrapped standard errors clustered at the judge level in parentheses. $^{*}p < .10$, $^{**}p < .05$, $^{***}p < .01$.

7.5 Treatment-effect parameters

Table 2 reports the key treatment effect parameters derived from the estimated marginal treatment effect curve. Several findings emerge.

First, the average treatment effect (ATE) is positive and statistically significant, indicating that incarceration increases subsequent offending on average. This masks substantial heterogeneity across individuals, as evidenced by the difference between the ATT and ATUT. The effect is larger for individuals who receive custody (ATT), while the effect for untreated individuals (ATUT) is smaller and statistically insignificant.

Second, the local average treatment effect (LATE) is positive and close in magnitude to the ATE, reflecting the effect for offenders at the judge-induced margin of treatment. This confirms that the instrumental variation identifies a policy-relevant subset of offenders.

Table 2: Treatment-effect parameters and interpretation (binary outcome; percentage points)

Parameter	Estimate	SE	p-value	Population	Interpretation	Per 1,000
LATE	0.199	0.024	< 0.001	Compliers	Custody increases recidivism by 19.9 pp for defendants whose custody status is shifted by judge assignment.	199
ATT	0.119	0.048	0.014	Treated ($D = 1$)	Among those sentenced to custody, custody increases recidivism by 11.9 pp relative to the counterfactual noncustodial sanction.	119
ATE	0.068	0.064	0.290	Population	Population-average effect; interpretation depends on overlap and extrapolation.	68
ATUT	0.060	0.077	0.436	Untreated ($D = 0$)	Effect of imposing custody on those typically noncustodial; estimate is positive but imprecise.	60

Notes: "Per 1,000" equals 1000× the percentage-point effect and should be read as additional reconvictions per 1,000 defendants in the relevant population (e.g., per 1,000 compliers for LATE). Standard errors follow the judge-cluster bootstrap.

Implications for marginal offenders. While the MTE estimates characterize how treatment effects vary with unobserved resistance, they do not directly reveal the observable characteristics of marginal offenders. However, the combination of a strong first stage and significant essential heterogeneity implies that marginal offenders differ systematically from both always-takers and never-takers. In particular, they are likely to be individuals whose incarceration status is sensitive to judge stringency—those near the custody threshold. Identifying their observable characteristics requires additional analysis that maps the estimated propensity score or latent resistance to observed covariates.

[show characteristics of marignal offenders here]

7.6 Hypothesis tests and structural interpretation

Table 3: Hypothesis tests

Hypothesis	Test	Contrast	Estimate	SE	p-value
H1	Sentencing-margin sign	LATE < 0	0.199	0.024	1.000
H2	Marginal vs treated	LATE $-$ ATT > 0	0.089	0.047	0.0293
H3	Marginal vs untreated	LATE $-$ ATUT < 0	0.139	0.067	0.9819
H4	External validity gap	LATE $-$ ATE $\neq 0$	0.131	0.053	0.014
H5	Reform derivative	MPRTE > 0	*Section 8*	*Section 8*	*Section 8*

The hypothesis results map directly to the structural model. H1 indicates that the IV-weighted average of the structural response function is positive, i.e., the relevant conditional average of $\sum_s w_{is} f(\eta_{is}(0))\Delta_{is}$ is positive for compliers. H2 (LATE>ATT) is consistent with an upward-sloping MTE(u) on the IV margin, implying larger net responses for higher-resistance types than for infra-marginal treated types. H4 confirms that uniform averaging (ATE) obscures the heterogeneity emphasized by the IV margin.

8 Policy simulations: PRTE and MPRTE

H5 needs policy simulations to estimate the PRTE and MPRTE. This section evaluates sentencing reforms using the estimated marginal treatment effect function MTE(u). The objective is to quantify the behavioral consequences of changing the fraction of defendants receiving custody. We focus on policy changes affecting the extensive margin (custody versus non-custodial sanction), not sentence duration.

8.1 Policy framework

Let treatment assignment follow a threshold rule

$$D(p) = \mathbf{1}\{U_i \leq p\},$$

where $U_i \sim \text{Unif}(0,1)$ is latent resistance to custody and p denotes the fraction of defendants sentenced to custody.

Consider a reform changing custody use from p_0 to p_1. Define the change in incarceration incidence as

$$\Delta D \equiv \mathbb{E}[D(p_1)] - \mathbb{E}[D(p_0)] = p_1 - p_0.$$

The resulting change in expected reoffending is

$$\Delta Y \equiv \mathbb{E}[Y(p_1)] - \mathbb{E}[Y(p_0)] = \int_{p_0}^{p_1} \text{MTE}(u)\, du.$$

We define the policy-relevant treatment effect (PRTE) as

$$\text{PRTE} = \frac{\Delta Y}{\Delta D}.$$

For infinitesimal reforms, letting $p_1 = p_0 + d\theta$, the marginal policy-relevant treatment effect (MPRTE) equals

$$\text{MPRTE} = \left.\frac{d\mathbb{E}[Y(p)]}{d\mathbb{E}[D(p)]}\right|_{p=p_0} = \text{MTE}(p_0).$$

Thus the MTE directly represents the causal effect of marginal sentencing reforms.

8.2 Identified policy domain

Based on the overlap of propensity scores and the width of confidence intervals in the MTE curve (see Figures 2 and 3, respectively), the empirically credible region of resistance is

$$u \in [0.05,\ 0.55].$$

Policy simulations are therefore conducted at $p_0 = 0.30$, a point well inside the region of common support.

8.3 Discrete reforms

We evaluate a plausible sentencing reform that changes custody use by five percentage points.

$$p_1 = p_0 \pm 0.05.$$

The resulting policy effects are

Table 4: Discrete policy reforms (PRTE)

Reform	ΔD	ΔY	PRTE
Increase custody ($0.30 \to 0.35$)	0.05	0.00410	0.08197
Decrease custody ($0.30 \to 0.25$)	-0.05	-0.00404	0.08074

The effects are symmetric and positive: expanding custody increases reoffending while reducing custody decreases reoffending.

8.4 Marginal reform

The marginal reform effect equals the MTE evaluated at the policy margin:

$$\text{MPRTE} = \text{MTE}(0.30) = 0.08127.$$

A one-percentage-point increase in custody therefore changes expected reoffending by

$$\Delta Y \approx 0.000813.$$

Interpreting $PRTE \approx MPRTE$. The similarity between the estimated PRTE and MPRTE is informative about the structure of treatment effect heterogeneity. The PRTE measures the impact of a finite reform that shifts the treatment probability for a non-negligible set of offenders, while the MPRTE measures the derivative of outcomes with respect to an infinitesimal policy change at the margin. If treatment effects vary sharply across the resistance distribution, these two quantities generally differ because discrete reforms reweight inframarginal individuals. In contrast, the near equality of PRTE and MPRTE implies that the marginal treatment effect curve is locally approximately linear over the region of policy-relevant support. In the context of the rational choice framework, this indicates that the net behavioral response to sentencing among marginal offenders changes smoothly with their latent resistance to punishment. Substantively, reforms that slightly expand the use of custody and reforms that moderately expand custody affect nearly the same population of offenders and therefore produce similar average behavioral responses. The empirical implication is that policy effects are driven by marginal offenders rather than large compositional shifts among inframarginal offenders.

8.5 Testing Hypothesis 5

These policy simulations support Hypothesis 5:

$$\text{MPRTE} > 0.$$

The positive MPRTE implies that marginal defendants – those whose incarceration status changes when sentencing policy shifts – experience a net increase in reoffending when exposed to custody. Because marginal policy changes move individuals precisely along this part of the resistance distribution, the relevant policy parameter is the MPRTE rather than the population ATE. For a clarification, we refer to Appendix B.

In this setting, changes in sentencing practice operate through selection into custody rather than through a uniform behavioral response across defendants. The MTE framework makes this distinction explicit by linking policy changes directly to the resistance distribution.

Note: PRTE and MPRTE are point estimates for now. We can consider producing bootstrapped SE's.

9 An optimal custodial policy rule

This section uses the estimated marginal treatment effect function to derive an optimal custody rate under a transparent policy objective. The purpose is not to recover a full structural model of the justice system, but to translate the estimated $\widehat{\text{MTE}}(u)$ into a normative decision rule on the empirically identified policy domain. Moreover, we show that $\widehat{\text{MTE}}(u)$ functions as sufficient-statistic for optimal custodial policy.

9.1 Policy lever and assumptions

We model custody assignment as a threshold rule,

$$D(p) = \mathbf{1}\{U_i \leq p\},$$

where $U_i \sim \text{Unif}(0,1)$ indexes latent resistance to custody and $p \in [0,1]$ is the policy-determined custody share. This representation is standard in the MTE framework and makes explicit that policy changes operate by moving a custody threshold through the resistance distribution.

Let $Y(p)$ denote the recidivism outcome induced by policy p. The key empirical input is the marginal treatment effect function

$$\text{MTE}(u) = \mathbb{E}[Y_i(1) - Y_i(0) \mid U_i = u],$$

which is interpreted throughout as a net behavioral response parameter. Under the threshold rule, expected outcomes satisfy

$$\mathbb{E}[Y(p)] = \mathbb{E}[Y(0)] + \int_0^p \text{MTE}(u)\, du. \tag{22}$$

We consider a policymaker minimizing a convex combination of expected reoffending and a resource cost of custody. Let $k \in [0, 0.10]$ be the per-custody cost expressed in "recidivism-equivalent" units. Up to a constant, the objective is

$$W(p) \equiv k\, p + \int_0^p \text{MTE}(u)\, du, \tag{23}$$

where the first term reflects that higher custody shares require more custodial resources, and the second term is the induced change in expected reoffending implied by the estimated MTE curve. Note that we do not consider the benefits of incarceration through incapacitation here. In Appendix G.2 we show that our conclusions of the paper do not depend on the inclusion or exclusion of incapacitation effects.

The policymaker minimizes the objective function $W(p)$ by choosing p:

$$min_p W(p) \tag{24}$$

Interpreting $W(\cdot)$ The objective function $W(p) = k\,p + \int_0^p \text{MTE}(u)\, du$ is written in *outcome units*. Because Y is an indicator (or probability) of reoffending over the follow-up window, both $\text{MTE}(u)$ and $\int_0^p \text{MTE}(u)\, du$ are measured in *probability points*. A euro-valued welfare objective can be obtained by multiplying $\mathbb{E}[Y(p)]$ by a monetary cost per recidivism event c_Y and multiplying p by the incremental monetary cost of custody c_D, in which case $k = c_D/c_Y$ is a dimensionless cost ratio. We perform such anlayses for the specific case of the Netherlands in Appendix G.1 below.

Interpreting k. The parameter k in (23) is a *shadow price* that converts a one-unit increase in custody prevalence into "recidivism-equivalent" units. Because the outcome Y is measured

as a probability (e.g., reconviction within a follow-up window), $k \in [0, 0.10]$ does *not* mean that custody is cheap in euro terms; it means that, under the chosen normalization, the policymaker is willing to accept at most 0.10 additional recidivism probability points per additional unit of custody (i.e., per one additional person treated in a unit mass population).

Equivalently, the marginal condition in Equation 25 below shows that k is the amount of increased reoffending probability the policymaker would tolerate to raise custody by one unit. If one instead expresses the objective in euros, k becomes the *ratio of a monetary custody cost to the monetary social cost of crime induced by recidivism.* In that case, even a very expensive custodial system can correspond to a numerically small k if the social cost of an additional reconviction is also large. We show this more explicitly in Appendix G.1

9.2 Analytical solution and comparative statics

Differentiating (23) yields

$$W'(p) = k + \text{MTE}(p), \tag{25}$$

and (when MTE is differentiable) the second derivative is $W''(p) = \text{MTE}'(p)$. An interior optimum $p^\star$ satisfies the first-order condition

$$\text{MTE}(p^\star) = -k. \tag{26}$$

Equation (26) shows that the optimal custody share equates the marginal effect of expanding custody, $\text{MTE}(p)$, with the (negative) marginal resource cost $-k$. When $\text{MTE}(p)$ is positive, increasing custody raises expected reoffending at the margin; in that case, even if custody resources were free ($k = 0$), the reoffending-minimizing policy would be to reduce custody on the identified domain. This result is also important for defining $\text{MTE}(u)$ as a sufficient-statistic as we will show in Section 9.4 below.

Lemma 2 (Corner solution under positive marginal effects)**.** *Suppose* $k \geq 0$ *and* $\text{MTE}(p) \geq 0$ *for all* p *in the feasible policy set* $\mathcal{P} \subseteq [0, 1]$*. Then* $W'(p) \geq 0$ *on* $\mathcal{P}$ *and the objective* (23) *is minimized at the smallest feasible custody share* $p^{\min} = \inf \mathcal{P}$*.*

Lemma 2 implies that when marginal custody increases reoffending throughout the feasible region, the optimal policy is a corner solution: choose the minimum feasible custody share.

Identified policy domain. Because the MTE is credibly identified only on the region of propensity-score overlap, we restrict policy evaluation to $\mathcal{P} = [p_L, p_U]$ where $p_L = 0.05$ and $p_U = 0.55$ (based on common support and the widening confidence bands near the tails). Policy recommendations outside $[p_L, p_U]$ would rely primarily on extrapolation rather than on the empirically identified MTE.

9.3 Numerical illustration

We implement (23) by interpolating the estimated MTE curve and approximating $\int_0^p \text{MTE}(u)\,du$ using trapezoidal quadrature on a fine grid.[3] We then minimize $W(p)$ over $p \in [p_L, p_U]$ for a range of custody-cost parameters k. Figure 4 plots the objective for a representative k, and Table 5 summarizes the argmin.

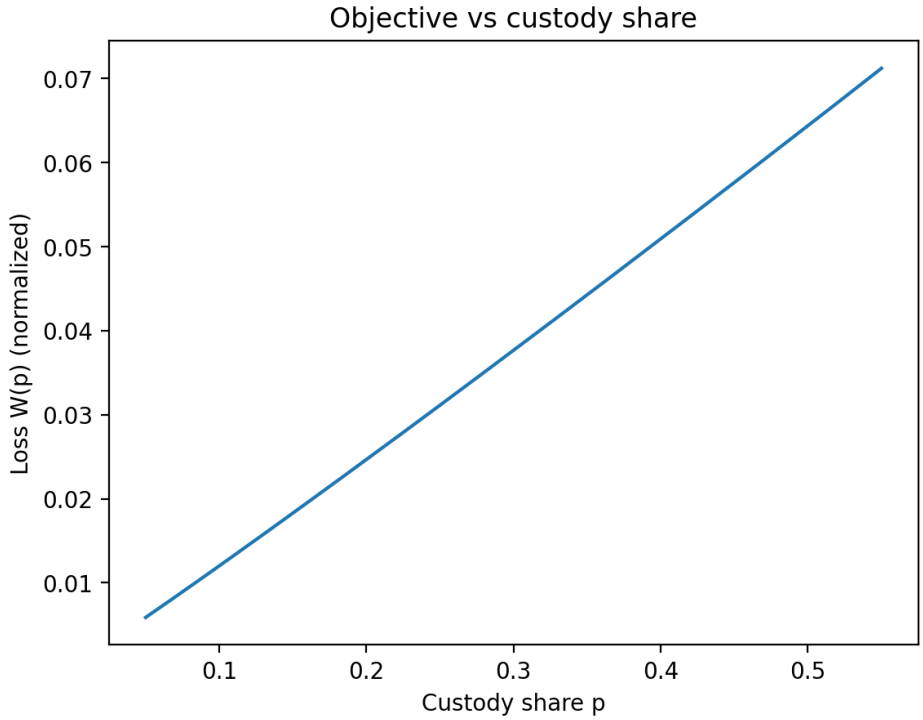


Figure 4: Objective $W(p) = kp + \int_0^p \widehat{\text{MTE}}(u)\,du$ over the identified domain $[0.05, 0.55]$ for $k = 0.05$.

Table 5: Optimal custody share on the identified domain using the estimated MTE curve

k	$u \in [p_L, p_U]$	$p^\star$	$\widehat{\text{MTE}}(p^\star)$	Policy
0.00	[0.05, 0.55]	0.05	0.0706	Minimize custody on u
0.02	[0.05, 0.55]	0.05	0.0706	Minimize custody on u
0.05	[0.05, 0.55]	0.05	0.0706	Minimize custody on u
0.10	[0.05, 0.55]	0.05	0.0706	Minimize custody on u

Interpretation. Although Figure 4 appears nearly linear, this is itself informative. Across reasonable nonnegative values of k, the optimum on the identified domain is $p^\star = p_L$. This

[3]The numerical simulation requires evaluating integrals of the estimated MTE curve, such as the PRTE, the welfare objective, and the policy frontier. These objects generally do not admit closed-form expressions because the estimated MTE is obtained semi-parametrically and need not follow a simple polynomial form over the identified support. Trapezoidal quadrature on a fine grid therefore provides a transparent and numerically stable approximation to the relevant areas under the MTE curve. The use of a fine grid is important because even a visually smooth MTE can exhibit local slope changes that matter for integrated policy objects. The procedure ensures that the numerical policy calculations reflect the estimated shape of the MTE rather than an additional approximation imposed.

result follows directly from the empirical finding that $\widehat{\text{MTE}}(u)$ is positive on the region where it is most credibly identified. In terms of the policy derivative, MPRTE = MTE(p), this means that marginal expansions of custody increase expected reoffending on the sentencing margin. The optimal-policy result therefore provides a normative restatement of Hypothesis 5: if MPRTE > 0 on the relevant domain, the reoffending-minimizing custody policy is to reduce custody to the lowest feasible level within that domain.

Why a corner-solution? A natural question is why the optimal custody share on the identified domain is a corner solution rather than an interior optimum. The answer follows directly from the marginal condition $W'(p) = k + \text{MTE}(p)$. On the empirically identified region of common support, the estimated $\widehat{\text{MTE}}(u)$ is positive throughout the domain where inference is most credible. With $k \geq 0$, this implies $W'(p) > 0$ over $[p_L, p_U]$, so welfare loss increases with custody intensity on that domain and the argmin is $p^\star = p_L$. This is not a modeling artifact; it is the normative translation of the empirical finding that marginal expansions of custody increase recidivism. An interior optimum would require either (i) a region where MTE(p) < 0 (marginal deterrence dominating net criminogenic effects), or (ii) an additional explicitly modeled benefit of custody not captured by Y (e.g., a monetized incapacitation term or a retributive value). We abstract from this to keep our model and conclusions simple and tractable.

Scope and limitations. The analysis is deliberately restricted to the overlap region $[p_L, p_U]$. Outside this region, p-counterfactuals would depend heavily on functional-form extrapolation of MTE(u). Moreover, the objective (23) is intentionally parsimonious: it treats MTE(u) as the empirically identified net effect of custody on reoffending and adds an explicit resource cost. Extensions that incorporate additional social components can be layered on top of the same marginal condition (26), but would require additional assumptions and data.

9.4 MPRTE as a sufficient-statistic

We now characterize which features of the treatment effect distribution determine the optimal custody threshold. The optimal policy p^* is defined by the first-order condition in Equation (26). That condition equates the marginal social benefit of custody to its marginal behavioral cost.

Proposition 1 (Local sufficiency of the marginal treatment effect)**.** *Let p^* denote an interior solution to Equation* (26)*. Then the optimal custody threshold depends on treatment effects only through the marginal treatment effect evaluated at p^*.*

For example, consider two judicial systems with identical cost parameters and identical values of MTE(p^*) *but different functions* MTE(u) *elsewhere on* $[0, 1]$*. Both systems imply the same optimal custody threshold.*

Proof. The welfare objective depends on treatment effects through the integral of MTE(u) over treated individuals. However, optimality requires the derivative of the welfare function with respect to the policy parameter p to vanish at p^*. By Equation (26), this derivative depends only on the value of MTE(u) evaluated at $u = p^*$. Changes in treatment effects for

inframarginal or extramarginal individuals therefore do not alter the optimality condition and leave p^* unchanged. □

Economic interpretation. The policymaker chooses custody prevalence by balancing marginal costs and benefits. The individuals whose treatment status does not change when p changes slightly are irrelevant for this trade-off. Only the behavioral response of offenders at the threshold matters.

Consequently, average treatment effects such as the ATE or ATT generally do not determine optimal policy, because they aggregate effects for individuals whose treatment status is fixed at the optimum. Instead, the relevant object is the treatment effect for offenders whose treatment status changes at the policy margin. This object is the marginal policy relevant treatment effect (MPRTE).

This result establishes a sufficient-statistics characterization of optimal sentencing policy: once the marginal response $\text{MTE}(p^*)$ is known, the remainder of the treatment effect distribution is not required to determine the optimal custody threshold.

10 Optimal custodial policy under uncertainty

Policymakers rarely know ex ante which individuals will be affected by a reform. In sentencing, a guideline change or enforcement directive may shift decisions unevenly across judges, courts, and case types. This section formalizes policy evaluation when the *reform margin* is uncertain and shows why the MPRTE—equivalently, the MTE evaluated at the status-quo threshold—is a sufficient statistic for the *first-order* (directional) effect of a small reform, but not generally for its magnitude.

10.1 Unknown reform margins and Bayesian beliefs

Let treatment be assigned by a threshold rule

$$D(p) = \mathbf{1}\{U \le p\}, \qquad U \sim \text{Unif}(0,1),$$

where U indexes resistance to custody. The marginal treatment effect is

$$\text{MTE}(u) = \mathbb{E}[Y(1) - Y(0) \mid U = u].$$

A reform of intensity θ changes the custody share from p_0 to $p(\theta)$, but the policymaker does not know ex ante which resistance types are shifted. We represent this uncertainty by a random *weight function* (kernel) $\omega_\theta(u)$ satisfying

$$\int_0^1 \omega_\theta(u)\, du = \Delta D(\theta) \equiv p(\theta) - p_0.$$

The policymaker's Bayesian beliefs are summarized by a prior distribution Φ over feasible kernels:

$$\omega_\theta \sim \Phi.$$

Examples of how Φ can be interpreted:

- **Sentencing guideline reform:** uncertainty reflects heterogeneous judicial reactions.
- **Enforcement intensification:** uncertainty reflects heterogeneous regional compliance.
- **Eligibility rule change:** uncertainty reflects heterogeneous take-up and administrative frictions.

Conditional on a realized kernel ω_θ, the induced change in expected outcomes is

$$\Delta Y(\theta \mid \omega_\theta) = \int_0^1 \text{MTE}(u)\, \omega_\theta(u)\, du,$$

and the ex ante expected policy effect is

$$\mathbb{E}_\Phi[\Delta Y(\theta)] = \int \left(\int_0^1 \text{MTE}(u)\omega(u)\, du \right) d\Phi(\omega). \tag{27}$$

10.2 Local reforms

We focus on *marginal* reforms. Formally, a reform is marginal if its intensity θ is small and therefore induces only a small change in treatment prevalence: $\Delta D(\theta) = O(\theta)$. The key economic idea is that small reforms shift treatment near the status-quo margin p_0. We encode this as follows.

Assumption 4 (Local margin uncertainty)**.** *For small θ, the realized kernel can be written as*

$$\omega_\theta(u) = \Delta D(\theta)\, k_\theta(u), \qquad \int_0^1 k_\theta(u)\, du = 1,$$

where the conditional distribution $k_\theta(u)$ has mean m_θ and variance v_θ satisfying

$$m_\theta = p_0 + O(\theta), \qquad v_\theta = O(\theta^2).$$

Assumption 4 allows policymakers to be uncertain about who is affected while ensuring that, for small reforms, the affected resistance types concentrate locally around p_0. The Taylor expansion allows us to turn this local concentration into an explicit approximation: it lets us express the average of MTE(u) over a distribution concentrated near p_0 in terms of the *level, slope, and curvature* of MTE(u) at p_0. This is necessary to assess the MTE(u) as a sufficient-statistic under policy uncertainty.

10.3 MPRTE as a first-order sufficient statistic under uncertainty

Proposition 2 (First-order sufficiency of the MPRTE)**.** *Suppose* MTE(u) *is twice continuously differentiable in a neighborhood of p_0 and Assumption 4 holds. Then for sufficiently small θ,*

$$\mathbb{E}_\Phi[\Delta Y(\theta)] = \mathbb{E}_\Phi[\Delta D(\theta)] \cdot \text{MTE}(p_0) + O(\theta^2).$$

In particular, if the policymaker expects the reform to expand custody on average (i.e., $\mathbb{E}_\Phi[\Delta D(\theta)] > 0$), the sign *of the expected reform effect is determined to first order by* $\text{MTE}(p_0) = \text{MPRTE}(p_0)$.

Proof. Write the realized reform effect as

$$\Delta Y(\theta) = \int_0^1 \text{MTE}(u)\omega_\theta(u)\,du = \Delta D(\theta)\int_0^1 \text{MTE}(u)k_\theta(u)\,du.$$

Perform a second-order Taylor expansion of $\text{MTE}(u)$ around p_0:

$$\text{MTE}(u) = \text{MTE}(p_0) + \text{MTE}'(p_0)(u-p_0) + \frac{1}{2}\text{MTE}''(p_0)(u-p_0)^2 + R_3(u),$$

where the remainder satisfies $|R_3(u)| \le C|u-p_0|^3$ locally. Integrating against $k_\theta(u)$ yields

$$\int \text{MTE}(u)k_\theta(u)du = \text{MTE}(p_0) + \text{MTE}'(p_0)(m_\theta - p_0) + \frac{1}{2}\text{MTE}''(p_0)v_\theta + O(\mathbb{E}|u-p_0|^3).$$

By Assumption 4, $(m_\theta - p_0) = O(\theta)$ and $v_\theta = O(\theta^2)$, so

$$\int \text{MTE}(u)k_\theta(u)du = \text{MTE}(p_0) + O(\theta).$$

Multiplying by $\Delta D(\theta) = O(\theta)$ implies

$$\Delta Y(\theta) = \Delta D(\theta)\text{MTE}(p_0) + O(\theta^2).$$

Taking expectations with respect to the prior Φ gives

$$\mathbb{E}_\Phi[\Delta Y(\theta)] = \mathbb{E}_\Phi[\Delta D(\theta)\text{MTE}(p_0)] + \mathbb{E}_\Phi[O(\theta^2)] = \text{MTE}(p_0)\mathbb{E}_\Phi[\Delta D(\theta)] + O(\theta^2),$$

because $\text{MTE}(p_0)$ is non-random given the status quo and $\mathbb{E}_\Phi[O(\theta^2)] = O(\theta^2)$ provided the remainder term $\mathbb{E}_\Phi[\Delta D(\theta)]$ is uniformly bounded in Φ. □

10.4 MPRTE and higer-orders

The approximation underlying Proposition 2 can be written more explicitly as

$$\Delta Y(\theta) = \underbrace{\Delta D\,\text{MTE}(p_0)}_{\text{first order}} + \underbrace{\Delta D\,\text{MTE}'(p_0)(m_\theta - p_0)}_{\text{shift in margin}} + \underbrace{\tfrac{1}{2}\Delta D\,\text{MTE}''(p_0)v_\theta}_{\text{dispersion in affected types}} + O(\theta^3).$$

Each term has a direct economic meaning:

First-order term (direction). The term $\Delta D\,\text{MTE}(p_0)$ is first order because it scales as $O(\theta)$: the reform moves only a small mass $\Delta D(\theta) = O(\theta)$, and those marginal individuals have treatment effect approximately $\text{MTE}(p_0)$. This term determines the *direction* of the effect for small reforms. As just proven in the previous section, the MPRTE is a sufficient statistic for the *directional* policy prediction under uncertainty.

Shift in margin (sensitivity to where the reform bites). The term $\Delta D\,\text{MTE}'(p_0)(m_\theta - p_0)$ captures that the realized reform may shift treatment slightly above or below the status-quo margin. If the policy implementation systematically targets a slightly different margin (e.g.,

prosecutors direct tougher bargaining specifically for a subset of cases), then $m_\theta - p_0 \neq 0$ and the slope $\text{MTE}'(p_0)$ determines how sensitive the policy effect is to that shift. Hence the MPRTE is not a sufficient-statistic for the treatment shift.

Dispersion in affected types (heterogeneous compliance). The term $\frac{1}{2}\Delta D\,\text{MTE}''(p_0)v_\theta$ captures that even if the reform is centered at p_0 on average, it may affect a dispersed set of resistance types (large v_θ). Greater dispersion makes curvature $\text{MTE}''(p_0)$ matter: if treatment effects change rapidly away from p_0, then implementation heterogeneity affects the magnitude of the expected policy impact. Hence the MPRTE is not a sufficient-statistic for the magnitude.

Why MTE is sufficient for direction but not magnitude. Proposition 2 shows that for small reforms the leading term depends only on $\text{MTE}(p_0)$. However, policy magnitude depends on the higher-order terms because both the location (m_θ) and spread (v_θ) of the reform margin are uncertain and determined by institutional details. Therefore, the MPRTE is a sufficient statistic for the *sign* of marginal policy effects, while the *magnitude* additionally depends on the slope and curvature of the MTE and on features of the reform margin distribution induced by Φ.

10.5 Connection to the empirical $MTE(u)$

Our empirical estimates imply that $\widehat{\text{MTE}}(u)$ is positive over the region of common support, and in particular $\widehat{\text{MTE}}(p_0) > 0$ for policy-relevant thresholds p_0 within that domain. By Proposition 2, this implies that any marginal reform that is expected to increase custody prevalence on average ($\mathbb{E}_\Phi[\Delta D(\theta)] > 0$) will increase reoffending to first order, even if policymakers cannot predict exactly which offenders are shifted.

Moreover, the estimated MTE curve is only mildly sloped on the empirically supported domain. This implies that the "shift in margin" term is small unless implementation systematically targets a different region of the resistance distribution. Finally, the curvature of the estimated MTE curve determines how sensitive policy magnitudes are to dispersion in affected types: a locally close-to-linear MTE—as in our case—implies that uncertainty about dispersion has limited quantitative consequences, whereas strong curvature would imply that implementation heterogeneity materially alters predicted reform effects.

Taken together, the empirical level, slope, and curvature of $\widehat{\text{MTE}}(u)$ provide a disciplined basis for evaluating both (i) the direction of marginal sentencing reforms under uncertainty and (ii) the robustness of predicted reform magnitudes to plausible implementation variation. For (i) $\widehat{\text{MTE}}(u)$ is proven to be a sufficient-statistic but not for (ii).

11 Conclusion

This paper characterizes optimal policy in environments with heterogeneous treatment effects. We show that the marginal policy-relevant treatment effect (MPRTE) is a sufficient statistic for optimal policy: interventions should expand if and only if the marginal response is favorable.

This result provides a transparent and operational link between causal inference and normative policy design.

A key contribution is that the sufficient-statistics result remains valid under policy uncertainty. Even when policymakers do not know which individuals will be affected by a reform, the MPRTE continues to determine the direction of optimal policy changes. This robustness makes the framework particularly relevant for real-world policy environments, where implementation margins are imperfectly understood.

Our findings complement the sufficient-statistics approach to policy evaluation (Saez, 2001; Chetty, 2009; Hendren, 2016, 2020). While that literature identifies reduced-form objects that are sufficient for welfare analysis, our results show that in environments with heterogeneous treatment effects, the sufficient statistic for optimal policy can be derived from the structure of the marginal treatment effect. The MPRTE summarizes the behavioral response at the policy margin and therefore provides a structural sufficient statistic for policy design. This connection highlights how causal inference methods can deliver the objects required for sufficient-statistics welfare analysis.

We apply this framework to custodial sentencing using quasi-random assignment of cases to judges and marginal treatment effect methods. The estimated marginal response to incarceration is criminogenic, implying that expanding custody increases recidivism at the policy margin. As a result, the reoffending-minimizing policy is to reduce custody within the empirically identified domain.

More broadly, the paper shows how heterogeneous treatment effect methods can be used to derive simple and robust policy rules. By reducing optimal policy to a single estimable object—the MPRTE—we provide a framework that is both theoretically grounded and empirically implementable. This approach can be applied to a wide range of policy settings in which interventions affect heterogeneous individuals through selection. Such settings arise in many areas of economics, including education (e.g., admission or remediation thresholds), labor markets (e.g., training program eligibility), health policy (e.g., insurance coverage rules), and public finance (e.g., benefit eligibility cutoffs).

References

Aizer, A. and Doyle Jr, J. J. (2015). Juvenile incarceration, human capital, and future crime: Evidence from randomly assigned judges. *Quarterly Journal of Economics*, 130(2):759–803.

Andresen, M. E. (2018). Exploring marginal treatment effects: Flexible estimation using stata. *The Stata Journal*, 18(1):118–158.

Andrews, D. W. and Buchinsky, M. (2000). A three-step method for choosing the number of bootstrap repetitions. *Econometrica*, 68(1):23–51.

Augustine, D. et al. (2022). Diversion and deterrence: Evidence from randomly assigned judges. *American Economic Journal: Applied Economics*. Forthcoming.

Becker, G. S. (1968). Crime and punishment: An economic approach. *Journal of Political Economy*, 76(2):169–217.

Bhuller, M., Dahl, G. B., Løken, K. V., and Mogstad, M. (2020). Incarceration, recidivism, and employment. *Journal of Political Economy*, 128(4):1269–1324.

Cameron, A. C. and Miller, D. L. (2015). A practitioner's guide to cluster-robust inference. *Journal of Human Resources*, 50(2):317–372.

Carneiro, P., Heckman, J. J., and Vytlacil, E. J. (2010). Evaluating marginal policy changes and the average effect of treatment for individuals at the margin. *Econometrica*, 78(1):377–394.

Chetty, R. (2009). Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods. *Annual Review of Economics*.

Collinson, R., Humphries, J. E., and Mader, D. (2021). Evictions and poverty in the united states. *American Economic Review*, 111(3):912–947.

Di Tella, R. and Schargrodsky, E. (2013). Electronic monitoring and recidivism: Evidence from argentina. *Journal of Law and Economics*, 56(2):399–432.

Dobbie, W., Goldin, J., and Yang, C. S. (2018). The effects of pretrial detention on conviction, future crime, and employment: Evidence from randomly assigned judges. *American Economic Review*, 108(2):201–240.

Dobbie, W. and Song, J. (2015). Consumer bankruptcy and financial health. *Journal of Finance*, 70(3):849–890.

Florens, J.-P., Heckman, J., Meghir, C., and Vytlacil, E. (2008). Identification of treatment effects using control functions in models with continuous, endogenous treatment and heterogeneous effects. *Econometrica*, 76(5):1191–1209.

Green, D. P. and Winik, D. (2010). Does incarceration reduce crime? *Criminology*, 48(2):357–386.

Harding, D. J., Morenoff, J. D., Nguyen, A. P., and Bushway, S. D. (2018). Short- and long-term effects of imprisonment on future felony convictions and prison admissions. *American Journal of Sociology*, 124(1):49–94.

Heckman, J. J. and Vytlacil, E. J. (2005). Structural equations, treatment effects, and econometric policy evaluation. *Econometrica*, 73(3):669–738.

Hendren, N. (2016). The policy elasticity. *American Economic Review*, 106(3):659–703.

Hendren, N. (2020). Measuring economic efficiency using inverse-optimum weights. *Journal of Public Economics*.

Kling, J. R. (2006). Incarceration length, employment, and earnings. *American Economic Review*, 96(3):863–876.

Leslie, E. and Pope, N. (2017). The unintended impact of pretrial detention on case outcomes: Evidence from new york city arraignments. *American Economic Journal: Applied Economics*, 9(3):28–57.

Loeffler, C. E. (2013). Estimating the effects of imprisonment on recidivism using a natural experiment. *Criminology*, 51(2):327–362.

Maestas, N., Mullen, K. J., and Strand, A. (2013). Does disability insurance receipt discourage work? using examiner assignment to estimate causal effects. *American Economic Review*, 103(5):1797–1829.

Saez, E. (2001). Using elasticities to derive optimal income tax rates. *Review of Economic Studies*, 68(1):205–229.

Wermink, H., Been, J., Schuyt, P., van Wijck, P., and Blokland, A. (2025). The price of retribution: Evidence from the willingness to pay for short-term prison sentences compared to community service orders. *Journal of Experimental Criminology*, 21:219–250.

Wermink, H. T., Blokland, A. A. J., Been, J., Schuyt, P. M., Tollenaar, N., and Apel, R. (2024). Estimating effects of short-term imprisonment on crime using random judge assignments. *Justice Quarterly*, 41(3):317–346.

A MTE and sentencing literature

Paper	**Country/ Setting**	**Dependent variable(s)**	**Endogenous variable (instrumented)**	**MTE?**	**Main conclusion**
Aizer and Doyle Jr (2015)	USA (juvenile court)	Education, adult crime	Juvenile incarceration (judge)	No	Incarceration worsens long-run ou comes.
Augustine et al. (2022)	USA (San Francisco)	Later convictions	Diversion referral (judge)	No	Diversion reduces subsequent offen ing.
Bhuller et al. (2020)	Norway	Recidivism, employment	Incarceration (judge)	**Yes**	Estimates MTE(u) and integrat policy effects.
Collinson et al. (2021)	USA (housing court)	Eviction, downstream outcomes	Eviction order (judge)	No	Evictions have persistent adverse fects.
Di Tella and Schargrodsky (2013)	Argentina	Recidivism	Electronic monitoring vs. prison (judge)	No	Monitoring reduces recidivism rel tive to prison.
Dobbie and Song (2015)	USA (bankruptcy court)	Credit, labor outcomes	Bankruptcy protection (judge)	No	Debt relief affects financial and lab outcomes.
Dobbie et al. (2018)	USA (bail court)	Crime, employment	Pretrial detention (judge)	No	Detention affects employment a crime.
Green and Winik (2010)	USA (Washington, DC)	Recidivism	Prison vs. probation (judge)	No	Limited deterrent effect of incarcer tion.
Harding et al. (2018)	USA (criminal courts)	Future convictions	Imprisonment (judge)	No	Effects differ across horizons and fenders.
Kling (2006)	USA (federal)	Employment, earnings	Incarceration length (judge)	No	Little long-run labor-market effect.
Leslie and Pope (2017)	USA (bail court)	Conviction, recidivism	Pretrial detention (judge)	No	Detention raises conviction probab ity.
Loeffler (2013)	USA (Cook County)	Recidivism, employment	Imprisonment (judge)	No	Effects depend on offender margin.
Maestas et al. (2013)	USA (Social Security)	Employment, earnings	DI receipt (examiner)	No	Disability insurance reduces lab supply.
Wermink et al. (2024)	Netherlands	Recidivism	Prison vs. noncustodial (judge)	No	Effects depend on institutional ma gin.

Table 6: Published judge/examiner leniency designs.

B Why the ATE is not the relevant policy parameter

A common interpretation of treatment-effect estimates evaluates whether a policy "works on average," corresponding to the average treatment effect

$$\text{ATE} = \int_0^1 \text{MTE}(u)\, du.$$

However, sentencing reforms do not randomly assign custody across the population. Instead, reforms alter the probability of custody near existing sentencing thresholds. Formally, a policy changes the treatment rule from $D = \mathbf{1}\{U \leq p_0\}$ to $D = \mathbf{1}\{U \leq p_1\}$, moving individuals locally along the resistance distribution rather than reassigning treatment uniformly.

The resulting behavioral effect is

$$\mathbb{E}[Y(p_1)] - \mathbb{E}[Y(p_0)] = \int_{p_0}^{p_1} \text{MTE}(u)\, du,$$

which depends only on the portion of the MTE(u) curve affected by the reform. The relevant policy parameter is therefore

$$\text{MPRTE} = \text{MTE}(p_0)$$

for marginal reforms and

$$\text{PRTE} = \frac{1}{p_1 - p_0} \int_{p_0}^{p_1} \text{MTE}(u)\, du$$

for discrete reforms.

The ATE instead averages treatment effects over individuals whose treatment status is unaffected by the reform. When treatment effects are heterogeneous, the ATE therefore generally differs from the effect of a policy change. Judge assignment designs are particularly informative because they identify effects precisely at the sentencing margin, making the MPRTE the natural causal parameter for evaluating sentencing practice.

C Robustness: Minimum caseload

[Caseload of min{20,50}]

D Robustness: Effect for 3 and 5 years

[Horizon might matter]

E Robustness: Effects by first-offense

[Helps disentangling deterrence vs. criminogenic by showing consistent patterns.]

F Robustness: Effects by offense types

[Helps disentangling deterrence vs. criminogenic by showing consistent patterns.]

G Extensions of the optimal policy rule

G.1 Numerical k for the Netherlands

The objective in (23) is expressed in outcome units; to discipline magnitudes, we rescale it into euros. Let $c_Y > 0$ denote the social cost (in euros) of one additional recidivism event within the follow-up horizon and let $c_D > 0$ denote the incremental direct resource cost of imposing custody rather than a non-custodial sanction on the index case. The euro-valued objective can be written (up to an additive constant) as

$$\widetilde{W}(p) = c_Y \int_0^p \mathrm{MTE}(u)\, du + c_D\, p,$$

so dividing by c_Y yields (23) with the dimensionless cost ratio

$$k = \frac{c_D}{c_Y}.$$

This makes clear why values of k in the range 0 to 0.10 can appear "small" numerically while still corresponding to large euro costs: k is a *ratio*, not a price level. A large custody cost c_D can still imply $k < 0.10$ if the social harm associated with an additional recidivism event c_Y is also large.

Disciplining c_D using Dutch cost accounting. Direct sanctioning costs for the Netherlands (in the year of the index-offense, 2012) imply that imprisonment is substantially more expensive than community service (Wermink et al., 2025). Using reported one-off administrative costs and per-day costs, the incremental direct cost of custody relative to community service can be approximated as

$$c_D(L) = \Delta c_{\text{one-off}} + \Delta c_{\text{day}} \cdot L,$$

where $\Delta c_{\text{one-off}} = 66.97 - 34.75$ euros, $\Delta c_{\text{day}} = 259 - 72.56$ euros/day, and L denotes the expected custodial duration in days.[4] For a short custodial spell of three weeks ($L = 21$ days), $c_D \approx 3{,}947$ euros.

Disciplining c_Y. The appropriate value of c_Y depends on the composition and social harm of crimes captured by the recidivism measure. We therefore treat c_Y as a calibrated policy parameter and report sensitivity. A conservative approach is to anchor c_Y to published per-offense victimization cost estimates and interpret it as the expected social cost of an additional reconviction event.[5] In the calibration below, we consider $c_Y \in \{10{,}000, 25{,}000, 50{,}000\}$ euros as a representative range.

[4] These are direct costs only; indirect costs related to reintegration and social security would increase c_D.

[5] Because reconviction aggregates heterogeneous delicts, c_Y should be interpreted as an average social harm per marginal event. Sensitivity over plausible ranges is therefore essential.

Implication for the magnitude of k. Combining $c_D \approx 3{,}947$ with $c_Y \in \{10{,}000, 25{,}000, 50{,}000\}$ yields

$$k = \frac{c_D}{c_Y} \in [0.079,\ 0.395].$$

Thus, even for short sentences, an empirically grounded calibration can imply values of k *above* 0.10. This underscores that the earlier normalization $k \in [0, 0.10]$ was not a claim about custody being inexpensive; it was merely a convenient scale in outcome units. In euro terms, custody is costly, and the implied k is plausibly substantial. However, given an upward sloping MTE(u) curve and our objective policy function, it is always optimal for the policymaker to minimize the share of custodial sentencing for the domain of common support.

G.2 Incapacitation and custody expansions

Our baseline objective interprets MTE(u) as the empirically identified net post-release behavioral effect of custody. However, custody may also generate a contemporaneous incapacitation benefit not captured in post-release recidivism. We therefore extend the objective by introducing an incapacitation benefit $b \geq 0$ (euros per custodial sentence):

$$\widetilde{W}(p) = c_Y \int_0^p \text{MTE}(u)\, du + (c_D - b)\, p.$$

This is especially relevant for applications other than the Netherlands because, empirically, incapacitation effects have been found to be negligible in the Netherlands (REFERENMCES HILDE). The marginal welfare effect with incapacitation benefits is

$$\frac{d\widetilde{W}(p)}{dp} = c_Y \text{MTE}(p) + (c_D - b),$$

so an interior optimum satisfies

$$\text{MTE}(p^\star) = \frac{b - c_D}{c_Y}.$$

Hence, if MTE(p) > 0 on the relevant domain, an interior optimum requires

$$b > c_D,$$

i.e. the incapacitation benefit must exceed the direct resource cost of custody. Using Dutch sanctioning cost data as in Section G.1 above, the incremental direct cost of custody relative to community service is $c_D \approx 3{,}947$ euros.

Mapping incapacitation into primitives. Let λ denote the expected number of crimes prevented per day of custody for the marginal offender, and c_{crime} the social cost per crime. Then

$$b(\lambda, L, c_{\text{crime}}) = \lambda \cdot L \cdot c_{\text{crime}}.$$

Substituting this into the interior optimum yields the threshold condition for custody expansion to be welfare improving:

$$\lambda \geq \frac{c_D + c_Y \text{MTE}(p)}{L \cdot c_{\text{crime}}}. \tag{28}$$

Numerical illustration. Using $L = 21$ days, $c_D = 3{,}947$ euros, $c_Y = 25{,}000$ euros and the estimated marginal effect at the lower bound of support $\text{MTE}(0.05) \approx 0.0706$, the above threshold in Equation 28 condition implies

$$\lambda^\star \approx 0.136 \text{ crimes/day}.$$

Thus, custody expansion would only be welfare improving if incarceration prevents roughly one crime every 7-8 (1/0.136) days for marginal offenders. This magnitude provides a transparent benchmark: the sign of optimal policy depends not on the existence of incapacitation per se, but on whether incapacitation is sufficiently large to offset both the direct cost and the criminogenic post-release effect.

When is custody expansion optimal? Equation 28 defines a frontier in (c_Y, λ)–space separating policy regimes in which custody expansion is welfare improving from regimes in which it is not. Rearranging,

$$\text{MTE}(p) \leq \frac{b - c_D}{c_Y} \quad \Longleftrightarrow \quad \lambda \geq \frac{c_D + c_Y \text{MTE}(p)}{L \cdot c_{\text{crime}}}.$$

The right-hand side defines the minimum incapacitation intensity required for marginal custody expansion to be socially optimal. Because $\text{MTE}(p) > 0$ on the empirically identified domain, the frontier is upward sloping in c_Y: when crime is socially more costly, the criminogenic post-release effect becomes more important, and custody must prevent more crime during incarceration to remain optimal. The frontier separates three regions:

- **Region I (custody contraction optimal):** λ below the frontier. Post-release criminogenic effects dominate incapacitation and costs.
- **Region II (knife-edge):** λ equals the frontier. Custody changes are welfare neutral.
- **Region III (custody expansion optimal):** λ above the frontier. Incapacitation benefits dominate both direct costs and behavioral responses.

Policy implication. The key implication is that optimal sentencing policy is not determined by whether incarceration "works" on average. Instead, it depends on whether incapacitation among marginal offenders is sufficiently strong relative to both costs and post-release behavioral responses. The MTE curve provides exactly the object needed to evaluate this trade-off because it isolates the effect for offenders whose treatment status changes under policy reform.