

Causally Learning an Optimal Rework Policy

Oliver Schacht*

OLIVER.SCHACHT@UNI-HAMBURG.DE

Sven Klaassen*[‡]

Philipp Schwarz*[†]

Martin Spindler*[‡]

Daniel Grünbaum[†]

Sebastian Imhof[†]

**Faculty of Business Administration, University of Hamburg, Moorweidenstraße 18, 20148 Hamburg, Germany*

†ams OSRAM, Leibnizstraße 4, 93055 Regensburg, Germany

‡Economic AI, Nürnberger Straße 262 A, 93059 Regensburg, Germany

Editor: Thuc Le, Jiuyong Li, Robert Ness, Sofia Triantafillou, Shohei Shimizu, Peng Cui, Kun Kuang, Jian Pei, Fei Wang, Mattia Proserpi

Abstract

In manufacturing, rework refers to an optional step of a production process which aims to eliminate errors or remedy products that do not meet the desired quality standards. Reworking a production lot involves repeating a previous production stage with adjustments to ensure that the final product meets the required specifications. While offering the chance to improve the yield and thus increase the revenue of a production lot, a rework step also incurs additional costs. Additionally, the rework of parts that already meet the target specifications may damage them and decrease the yield. In this paper, we apply double/debiased machine learning (DML) to estimate the conditional treatment effect of a rework step during the color conversion process in opto-electronic semiconductor manufacturing on the final product yield. We utilize the implementation `DoubleML` to develop policies for the rework of components and estimate their value empirically. From our causal machine learning analysis we derive implications for the coating of monochromatic LEDs with conversion layers.

Keywords: Causal Inference, Machine Learning, Heterogeneous Treatment Effects, Double Machine Learning, Policy Learning

1 Introduction


In the face of an ever-increasing pressure to produce high quality products at low costs, manufacturing companies are striving to transform products that do not meet the quality targets into sellable units by reworking them (Liu et al. (2009)). For ecological reasons, too, reworking is preferable to discarding products that are not fulfilling preset target specifications.

The emergence of polychromatic light-emitting diodes (LEDs) extended the range of applications to areas in which traditional white light sources seemed predominant (e.g. general home lighting, streetlamps, automotive, LCD displays), gradually superseding them. Since then, multiple approaches for the production of white-emitting LEDs have been developed (for an in-depth introduction, see Cho et al. (2017)), with the most successful one being the coating of a monochromatic blue LED with a phosphor conversion layer. The coating layer partially shifts the emitted wavelength spectrum of the light source, resulting in a different perceived color (Figure 1a). During manufacturing this so-called conversion step is crucial in order to produce LEDs with the color properties that meet the customers requirements.

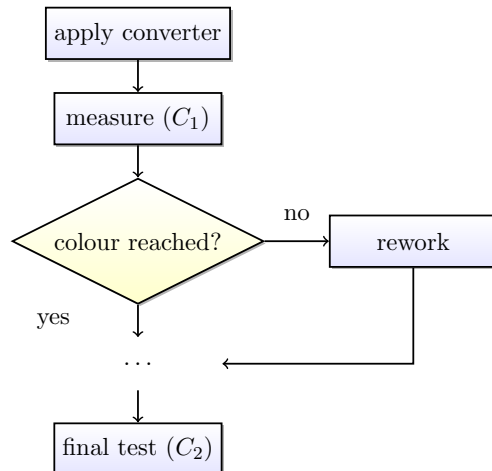
In order to ensure a high yield, that is a high share of LEDs in a production lot that match the target color specifications, the rework workflow shown in Figure 1b is used. Starting at a mean color point C_0 , multiple conversion layers are applied to the chips of a production lot, resulting in the converted mean color point C_1 . Using the measurement C_1 , a decision is made, whether the chips are coated with an additional conversion layer. After this optional rework step, production continues until the final products are assembled and tested. Finally, the product specification is checked for each LED in the production lot using the test results (see mean color point C_2 in Figure 1a). Note that there are steps in between rework and the final testing that have an influence on the final color point. That is a certain color shift needs to be taken into account when applying the conversion material.

This process setup poses the question for which color points C_1 the additional rework step should be applied in order to improve the yield of a production lot. The currently employed decision rule is based on the experience of the quality managers.

In the literature, rework policies are often determined by methods of planning and control. For an overview of methodologies, see Flapper et al. (2002). The recent extension of machine learning (ML) from purely predictive tasks to causal questions enables the use



(a) CIE 1931 color space chromaticity diagram



(b) Simplified production workflow

Figure 1: Conversion process including mean color point measurements before (C_0) and after conversion (C_1), as well as for the final LEDs (C_2) of a production lot.

of ML for learning decision policies. In this work, we are going to utilize double machine learning to derive rework policies based on real production data from the phosphorous conversion process of blue LEDs.

2 Setting

2.1 Data

The collected data consists of observations $W_i = \{Y_i, A_i, X_i\}$, where Y_i denotes the outcome, which is a measure of the final yield of each production lot. Further, $A_i \in \{0, 1\}$ represents the treatment/action variable, corresponding to the decision whether to rework a production lot, and X_i are characteristics of the production lot before assignment of a rework. Finally, we use data of two different product types ($P1$ and $P2$) with $n_{P1} = 32,669$ and $n_{P2} = 17,802$ observations. In the observed data, there is no clear policy in place. The quality managers assign the rework treatment by visually inspecting the plot of the covariates $X_i = \{(C_{1,x}^i), (C_{1,y}^i)\}$, consisting of the CIE coordinates of the mean color point C_1^i after the conversion. Thus, under the current policy, there is an overlap of production lots with similar characteristics after conversion but with different treatments assigned.

As indicated by the color points C_0 , C_1 , and C_2 in Figure 1a, the application of the conversion material shifts the resulting color point along a curve in the CIE color space from an unconverted monochromatic emission (C_0) to a fully converted one (C_2). Thereby, the amount of applied conversion material controls the position on this *conversion curve*. Minor deviations orthogonal to the conversion curve can be explained through process instabilities. Due to the scale of the C_1 measurements, the conversion curve can be well approximated by a linear function. Thus, the application of a principal component analysis (PCA) leads to transformed color coordinates $\text{PCA}(C_1) = (C_m, C_s)$ where the first principal component C_m captures the position on the conversion curve representing the main decision criteria, and the second principal component C_s is mainly determined by process fluctuations. In the following, we refer to C_m as the *main color point measure* and to C_s as the *secondary color point measure*.

Figure 2 represents the assumptions on the underlying causal structure in form of a directed acyclic graph (DAG, see Pearl (1995)). The yield of a production lot depends to a large extent on how close the chips are to the target color location in the final quality control. As the rework decision is based on the mean color point measurements after phosphorous conversion, the luminous properties of the opto-electronic chips influence not only




Figure 2: Causal Graph of the Rework Treatment Problem

the yield but also the decision to rework a specific product lot. Consequently, the luminous properties act as confounders with respect to the rework effect. Although there are other influencing factors on the treatment decision (i.e. the individual employee) or the outcome (i.e. variations in downstream process steps), these do not affect both simultaneously.

2.2 Effect Identification

As argued in Section 2.1, the rework decision is done based on the characteristics X . Consequently, conditioning on X will close all backdoor paths between rework decision and yield. Put differently, we base our identification on conditional exogeneity, such that

$$Y(a) \perp A | X,$$

where $Y(a)$ denotes the potential outcome for $a \in \{0, 1\}$ (see Rubin (2005)). Under the assumptions above, the underlying causal structure can be represented by an interactive regression model (IRM):

$$\begin{aligned} Y &= g_0(A, X) + U, & \mathbb{E}[U | X, A] &= 0, \\ A &= m_0(X) + V, & \mathbb{E}[V | X] &= 0, \end{aligned}$$

where the conditional expectations

$$g_0(A, X) = \mathbb{E}[Y | X, A], \tag{1}$$

$$m_0(X) = \mathbb{E}[A | X] = P(A = 1 | X) \tag{2}$$

are unknown and might be complex functions of X . In this structural equation model, the average treatment effect (ATE)

$$\theta_{ATE} = \mathbb{E}[g_0(1, X) - g_0(0, X)]$$

as well as the average treatment effect of the treated (ATTE)

$$\theta_{ATTE} = \mathbb{E}[g_0(1, X) - g_0(0, X) | A = 1]$$

are identified. The work of Chernozhukov et al. (2018) enables the use of machine learning (ML) algorithms such as random forest or boosting to obtain precise estimates of treatment effects combined with confidence intervals to access estimation uncertainty.

In double machine learning, inference is based on a method of moments estimator

$$E[\psi(W; \theta_0, \eta_0)] = 0, \tag{3}$$

with $\psi(W; \theta_0, \eta_0)$ being a Neyman-orthogonal score that identifies the causal parameter θ_0 , given ML estimates of the parameters or functions η_0 (referred to as nuisance). The Neyman-orthogonality property ensures robustness of the estimator against small perturbations in the ML nuisance estimates. To ensure that the perturbation is small enough, the nuisance estimators have to converge sufficiently fast (usually rates of $n^{-1/4}$ are required). For most ML learners, this indirectly implies smoothness assumptions on the true nuisance functions. In this application, it appears reasonable to assume smoothness as the nuisance function

in (1) corresponds to the physical process that describes the influence of adding more phosphorous coating to a chip. The nuisance in (2) is determined by individual decisions based on the color space, which suggests quite smooth probabilities. Further, the usage of k -fold crossfitting safeguards against overfitting and enables to control the complexity of the estimator, leading to appealing properties such as \sqrt{n} -consistency and approximate normality. For more details, we refer to Chernozhukov et al. (2018).

In case of an IRM, the score for an ATE estimator is given by the linear form

$$\psi(W_i; \theta, \eta) := \psi_a(W_i, \eta)\theta + \psi_b(W_i, \eta) \quad (4)$$

$$\begin{aligned} &= -\theta + g(1, X_i) - g(0, X_i) \\ &\quad + \frac{A_i(Y_i - g(1, X_i))}{m(X_i)} - \frac{(1 - A_i)(Y_i - g(0, X_i))}{1 - m(X_i)}, \end{aligned} \quad (5)$$

which is also known as augmented inverse propensity weighting (AIPW) (Robins and Rotnitzky (1995)).

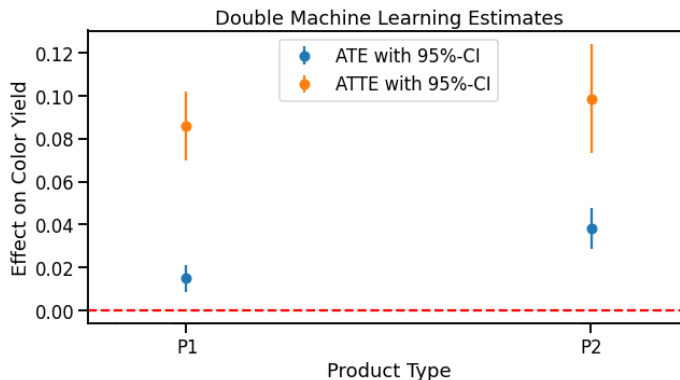


Figure 3: Average treatment effect (ATE) of reworking a chip lot and average treatment effect on the treated (ATTE) estimated by double machine learning.

Applied to the data, the estimates (see Figure 3) suggest that for both products, the treatment effect is much higher on the treated than on the whole population. This emphasizes that treating all individual lots is not desirable in this setting, as the effect varies over the covariates X .

Figure 4 shows the distribution of the main color point measure in X for the treated and non-treated subsamples. There is no clear cut-off due to the reasons mentioned above. Ideally, there would be a policy in place that assigns the treatment to each production lot if the corresponding effect is above a set threshold (incorporating costs etc.). Our goal is to estimate a policy $\pi \in \Pi$, that maps individuals features $X_i \in \mathcal{X}$ to a treatment decision: $\pi : \mathcal{X} \rightarrow \{0, 1\}$, which incorporates the heterogeneity of the effect of the rework.

During the rest of the main analysis, for simplicity, we focus on the product type P1 (results for P2 can be found in Appendix B).

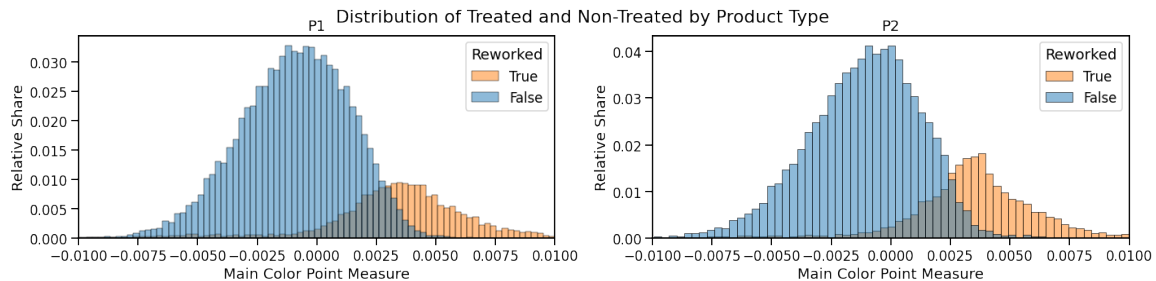


Figure 4: Distribution of the main color point measurement. Observations with a low value are already close to the target, thus they have a low treatment probability. Observations with a high value are therefore treated with a higher probability.

3 Methodology

As motivated in the previous section, in this scenario the treatment effect is heterogeneous with respect to covariates X . We will employ different approaches from conditional treatment effect estimation and optimal policy learning to learn a rework policy in this setting.

3.1 Conditional Average Treatment Effects

Given a set of the covariates \tilde{X} (not necessarily included in X), the conditional average treatment effect (CATE) is defined as

$$\theta_0(\tilde{x}) := \mathbb{E}[g_0(1, X) - g_0(0, X) | \tilde{X} = \tilde{x}].$$

Semenova and Chernozhukov (2021) propose to approximate $\theta_0(\tilde{x}) \approx b(\tilde{x})^T \beta$ via a linear form, where $b(\tilde{x})$ is a d -dimensional basis vector of \tilde{x} . The idea is based on projecting the part $\psi_b(W_i, \eta)$ of the Neyman-orthogonal scores onto the predefined basis vector $b(\tilde{x})$. The authors provide extensive theory to construct pointwise and jointly valid confidence intervals via gaussian approximations. As explained in Section 2.1, being close to the color target might be well approximated via the first principal component of the production lot properties X .

To access the heterogeneity associated with the first principal component, we set \tilde{X} to be the first principal component of the covariates X . Further, to allow for a flexible conditional treatment effect, we use a b-spline basis for $b(\tilde{x})$. The effect estimates combined with confidence intervals are displayed in Figure 6.

As a positive CATE indicates improvements in yield, our estimated policy will be based on whether the effect on the yield surpasses some given threshold

$$\hat{\theta}(\tilde{x}) \geq \gamma,$$

where \tilde{x} is the first principal component of X , and the threshold $\gamma > 0$ can be used to incorporate costs and select rework groups at different effect levels.

Additionally, considering the second principal component can provide further information about treatment effect heterogeneity and might improve the resulting policy. To be able to approximate the conditional effect flexibly based on the first two principal components, we rely on a tensor product of b-splines to construct our basis vector $b(\tilde{x})$.



Figure 5: Principal components of the chip properties before treatment, treatment decision and observed outcome in the data for product type “P1”.

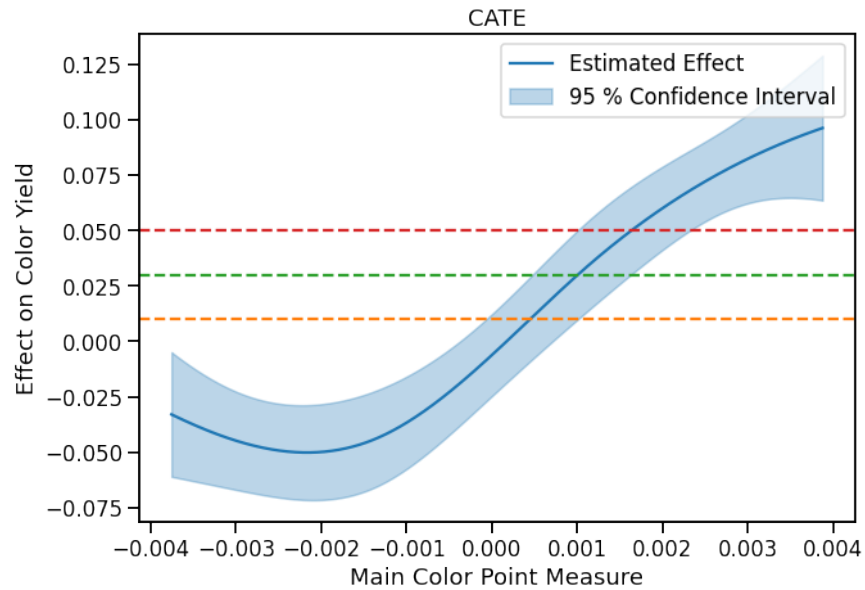


Figure 6: Estimation of the CATE along the first color point measurement. The horizontal dashed lines represent different policy thresholds γ .

The resulting effect estimates are displayed in Figure 7, where we omitted the confidence regions for clarity. It is visible that the first principal component captures the main heterogeneity on the treatment effect. Nevertheless, we will compare the difference in the policies in Section 4.




Figure 7: CATE with respect to both color point measurements.

In both cases above, instead of using the effect estimate, using the lower bound of the confidence interval will result in more conservative policies but account for the estimation uncertainty (see Appendix A.3). Additionally, one might consider the percentage of the reworked panels when deciding which threshold might be preferable to incorporate capacity constraints.

3.2 Policy Learning

An alternative approach is to directly estimate a policy based on the part $\psi_b(W_i, \eta)$ in Equation 5 as introduced by Athey and Wager (2021). They propose to estimate the treatment assignment rule as

$$\hat{\pi} = \arg \max_{\pi \in \Pi} \frac{1}{n} \sum_{i=1}^n (2\pi(X_i) - 1) \psi_b(W_i, \hat{\eta}). \quad (6)$$

Given a specified policy class Π and the corresponding regret

$$R(\pi) := \max_{\pi' \in \Pi} \{ \mathbb{E}[Y_i(\pi'(X_i))] \} - \mathbb{E}[Y_i(\pi(X_i))],$$

the authors are able to derive regret bounds of order $1/\sqrt{n}$.

The learning problem in Equation 6 can be reformulated as a weighted classification problem

$$\hat{\pi} = \arg \max_{\pi \in \Pi} \frac{1}{n} \sum_{i=1}^n \lambda_i H_i (2\pi(X_i) - 1), \quad (7)$$

with weights $\lambda_i = |\psi_b(W_i, \hat{\eta})|$ and target $H_i = \text{sign}(\psi_b(W_i, \hat{\eta}))$. Thus, we define the space of possible policies Π to be the class of depth- m decision trees (see Figure 8 as an example) with $m \in \{1, 2\}$. Appendix A includes results for Π being the class of linear decision boundaries after application of the radial basis function (rbf) kernel (using support vector machines). To evaluate policies at different thresholds, we reduce the score by the desired threshold $\psi_b(W_i, \hat{\eta}) - \gamma$ before classification.




Figure 8: Exemplary policy tree of depth 2.

4 Results

The implementation of the policy estimation relies on the `DoubleML`-package (Bach et al. (2022, 2021)). The package allows for direct estimation of the IRM and the CATEs with a variety of machine learning algorithms. The implemented `DoubleML` model is based on the original data without principal component transformation. As the quality of the policy estimation will improve with high-quality estimates of the unknown nuisance elements $\eta = (g_0, m_0)$ (see (1) and (2)), we rely on tuned machine learning algorithms. In the main analysis, we use a tuned `LightGBM` (Ke et al. (2017)) estimator. Tuning was performed by searching the hyperparameter space with `FLAML` (see Wang and Wu (2019)). Further, we employ 5-fold crossfitting and trim the propensity score estimate \hat{m} at 0.025 and 0.975, respectively. As basis vectors $b(\tilde{x})$ for the CATE estimation we construct cubic b-splines with 5 degrees of freedom in the one-dimensional case and a tensor product of quadratic b-splines with 5 degrees of freedom in the two-dimensional case. Appendix A contains evaluation and sensitivity checks with different learners and policies based on the confidence intervals.

Using the estimated score elements $\psi_b(W_i, \hat{\eta})$ from the `DoubleML`-package, we can apply weighted classifiers to obtain the policy learning results. As mentioned by Athey and Wager (2021), greedy classification trees will not achieve the optimal regret bound as they do not optimize properly over the whole policy space. Consequently, we rely on the R-package `policytree` (Sverdrup et al. (2020)) to estimate the policy with exact tree search (results using greedy trees are available in Appendix A.1).

To illustrate the effect of different thresholds, we consider $\gamma \in \{0.01, 0.03, 0.05\}$ and report the corresponding policy.

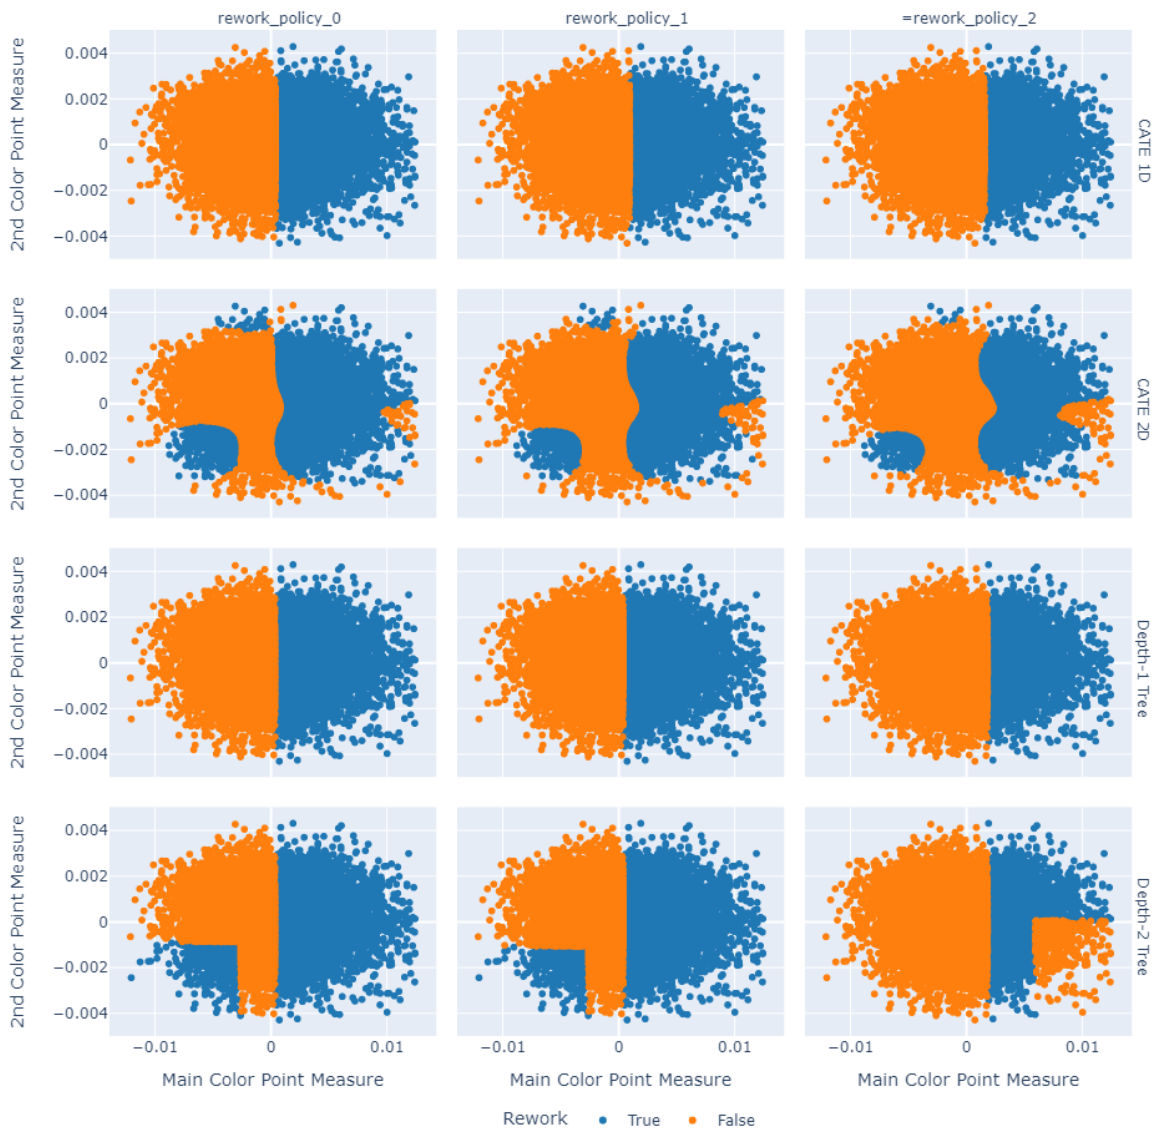


Figure 9: Comparison of the estimated policies. The rows contain the different policy learning approaches, while the columns refer to different cut-off levels γ .

As it is shown in Figure 9, all methodologies lead to a major split along the main color measurement. This is in accordance to the domain knowledge described in Section 1. While the one-dimensional CATE policy and the depth-1 tree are not flexible enough to display splits along the secondary color measurement, the two-dimensional CATE and the depth-2 tree split the area further. Nevertheless, all methodologies lead to similar decision rules.

Table 1 contains the share of observations that would have to be treated under the estimated policies. Even under the largest threshold (5%) there would be a significant increase in workload from reworking, which nearly doubles for $\gamma = 0.01$. This suggests that,

Method	Policy 0	Policy 1	Policy 2
CATE 1D	0.3864	0.3352	0.2806
CATE 2D	0.4179	0.3599	0.2973
Depth-1 Tree	0.4309	0.3842	0.2555
Depth-2 Tree	0.4708	0.4222	0.2393

Table 1: Share of reworked panels per policy estimate (Share of treated in the observations: 0.2101).

neglecting capacity restrictions, there could be improvements in overall yield by employing a data-driven rework policy.

To evaluate the average effect of the different policies, we estimate the group average treatment effect (GATE) for the reworked panels under each policy (ATTE for the observational data or current policy).

Method	Policy 0	Policy 1	Policy 2
CATE 1D	0.0748	0.0790	0.0880
CATE 2D	0.0717	0.0811	0.0867
Depth-1 Tree	0.0699	0.0755	0.0949
Depth-2 Tree	0.0686	0.0738	0.1026

Table 2: GATE by suggested treatment groups of the policy estimation (GATE in observed policy: 0.0903).

Surprisingly, the current observational policy still has a very high average effect compared to the estimated policies. Based on our analysis, the average treatment effect on the group of to-be-treated chip lots improves only for the 5%-threshold tree policies. However, a smaller average effect per panel is not necessarily disadvantageous, as the share of treated is increasing.

To quantify the potential improvement in overall yield with respect to the corresponding policy, we weight the GATE with the share of reworked panels.

Method	0	1	2
CATE 1D	0.0289	0.0265	0.0247
CATE 2D	0.0300	0.0292	0.0258
Depth-1 Tree	0.0301	0.0290	0.0242
Depth-2 Tree	0.0323	0.0312	0.0246

Table 3: Value of the estimated policies (Value of observed policy: 0.0178).

The results in Table 3 suggest a large improvement in policy value for all data-driven policies. The improvement is mainly driven by a larger fraction of reworked panels. Even if capacity constraints do not allow for a rework of 40%, e.g. the one-dimensional CATE suggests that 0.62% more yield could be achieved by reworking 6.62% additional chip lots.

5 Conclusion

In this paper, we have used causal machine learning to estimate different rework policies for the color conversion process in opto-electronic semiconductor manufacturing. The robustness and similarity of the different approaches support the original argument, that the first principal component of the C_1 measurement contains most of the relevant information for a rework decision. This highlights the idea that domain knowledge can be very helpful to design simple but effective data-driven policies. Consequently, to improve the current rework decision rule, we propose a policy where the treatment decision is based on the main color point measure C_m . This would increase the number of reworked panels and still select the panels where the rework is the most effective in increasing the overall yield. Policy learners such as the two-dimensional CATE and the depth-2 tree, by being more flexible, also suggest splits along the secondary axis. However, we are cautious about these as they might as well be a product of overfitting. As visible in Figure 5, the overlap of the data is mostly limited to the interval of $C_m \in [-0.005, 0.005]$. Outside this area, predictions overemphasize single data points, leading to overfitting. The confidence intervals in these regions are considerably large (including zero), which underlines the substantial amount of statistical uncertainty within these regions.

In conclusion, it is very important to combine implications from causal inference with the knowledge of domain experts to make value out of the data. Carefully weighting up the possible cut-offs and also taking into account how many reworks are feasible with production capacities, we suggest implementing a clear cut-off along the main measurement which allows an easy implementation and thus higher compliance of the decision makers. Furthermore, at a cut-off, the real effect could be then measured with techniques such as regression discontinuity to allow for further analysis.

Acknowledgments and Disclosure of Funding

This work was funded by the Bavarian Joint Research Program (BayVFP) – Digitization (Funding reference: DIK0294/01). The research partners ams Osram and Economic AI kindly thank the the VDI/VDE-IT Munich for the organization and the Free State of Bavaria for the financial support.

Appendix A. Robustness checks

A.1 Additional Policy Learning Results

In this section, we present additional results for policy learning. The policy estimation with greedy decision trees relies on the `DecisionTreeClassifier` from `scikit-learn` (Pedregosa et al. (2011)). We also include a weighted classification with a rbf kernel in a support vector machine. However, we caution that the theoretical bounds by Athey and Wager (2021) do not apply to this learning method, since they are based on bounds on the VC-dimension of the policy class Π .

Figure 10 contains the resulting policies for the tuned ML LGBM learner as in the main analysis. We observe that the greedy tree search algorithm finds no splits along the sec-

ondary color measurement. The rbf kernel encircles little areas of color point measurements with a higher value. These areas are very similar to the two-dimensional CATE. However, they contain only few observations, so the evidence from them is limited.

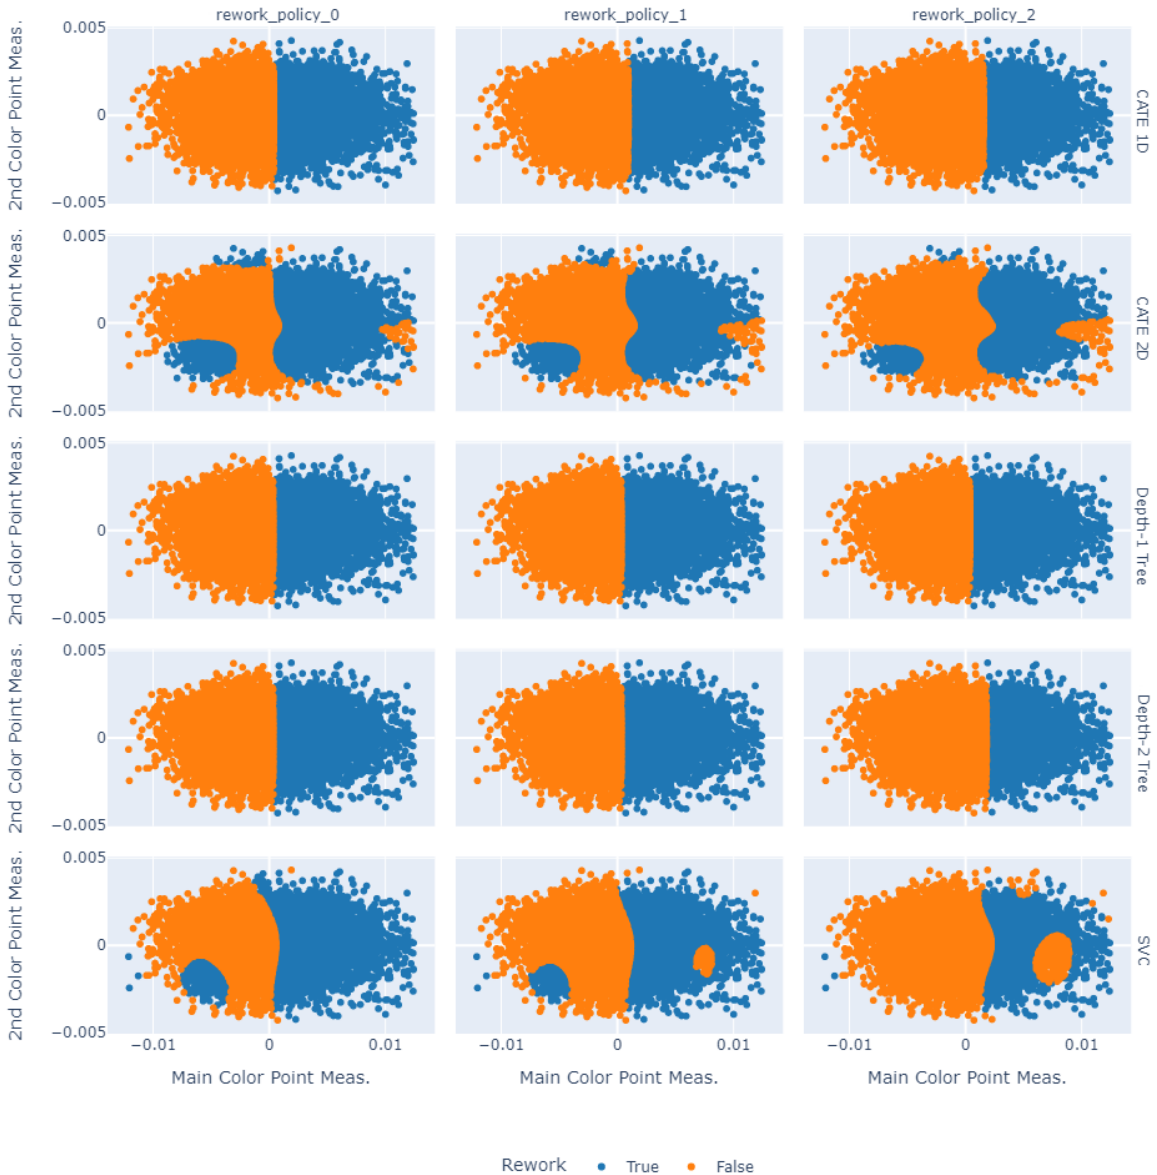


Figure 10: Comparisons of the estimated policies with additional methods.

A.2 DoubleML Machine Learner Comparisons

In this section, we explore the robustness of our estimates by varying the ML method which is used to construct the estimators in Equations 1 and 2. We compare the identical policy learning approaches as in Appendix A.1. The comparison includes results for XGBoost,

random forest and linear learners. Table 4 contains the *RMSE* of the respective ML methods in the nuisance estimation.

Learner	RMSE $m(x)$	RMSE $g_0(x)$	RMSE $g_1(x)$
Tuned LGBM	0.2821	0.8862	1.001
XGBoost	0.2857	0.9078	1.043
Random Forest	0.2983	0.9581	1.077
Linear Models	0.2925	0.9682	1.083

Table 4: RMSE for different nuisance learners.

Both gradient boosting frameworks achieve good predictive performance. Random forest and linear model show a slightly weaker performance. In Figures 11, 12 and 13 we show the results for the policies (to allow comparisons to the learner in the main analysis in Figure 10). The results seem to be rather robust across non-linear learners, especially in the main component split. Linear models achieve nearly the same performance, but, compared to the boosting learners, the more complex policy estimates seem less stable. This might indicate that linear projections are sufficient for simple policies, but might not capture all, potentially complex, dependencies in the data. As visible in Figure 5, in the center of the cloud of observations, an Euclidean plane likely yields low approximation errors on the outcome, given the covariates. Consequently, the simple policy learning approaches such as one-dimensional CATE or depth-one policy tree show similar results as with non-linear learners. However, moving away from the center of the points, a linear approximation might not be sufficient.




Figure 11: Comparison of the estimated policies with the nuisance components estimated by gradient boosting implementation `xgboost`.

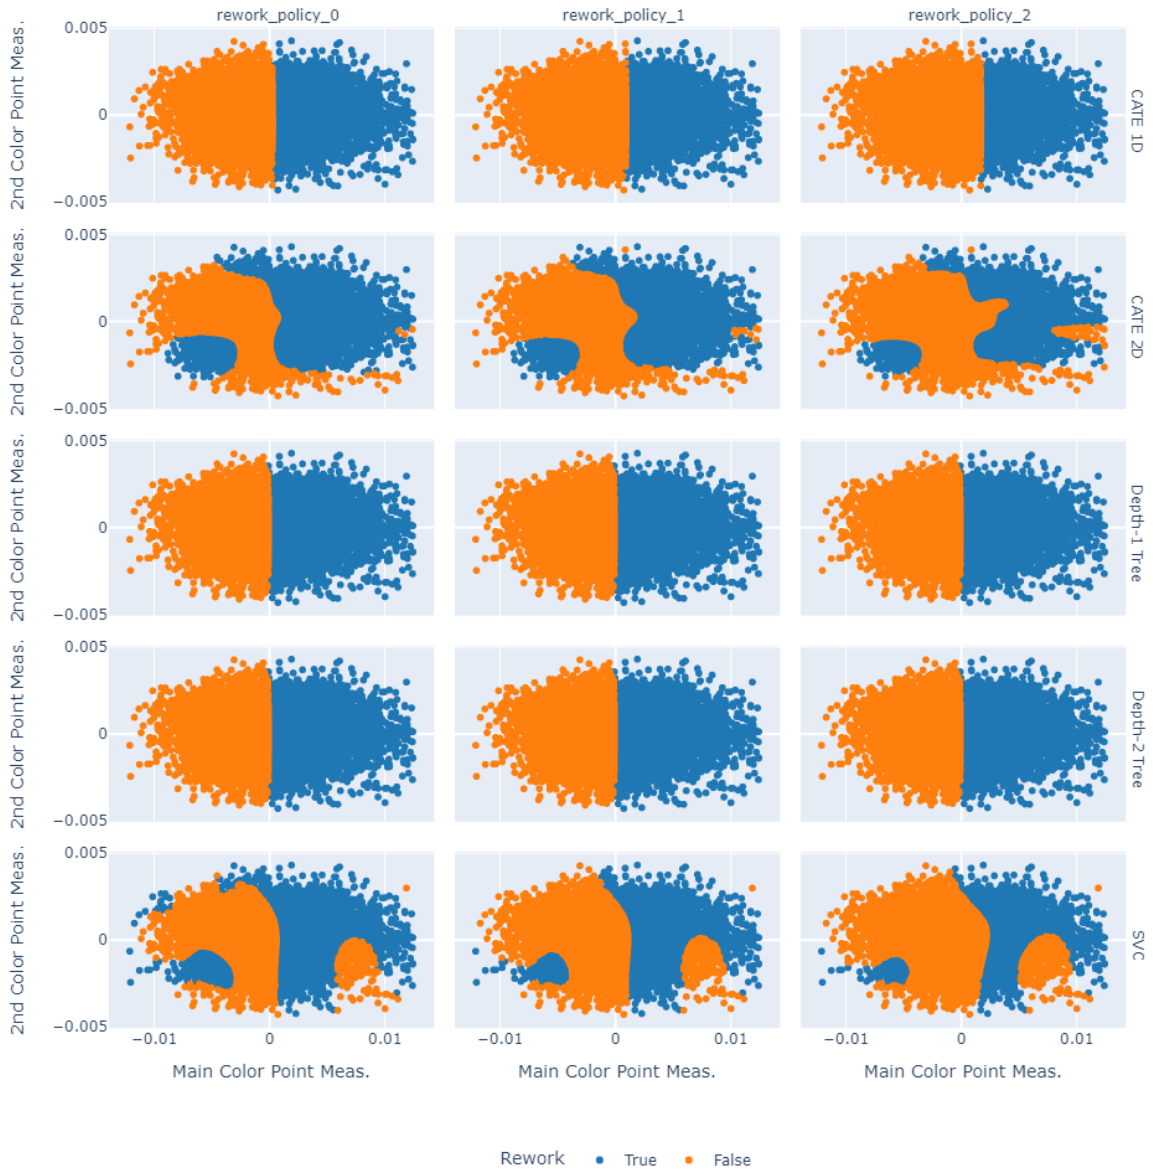


Figure 12: Comparison of the estimated policies with the nuisance components estimated by the random forest implementation of `scikit-learn`.

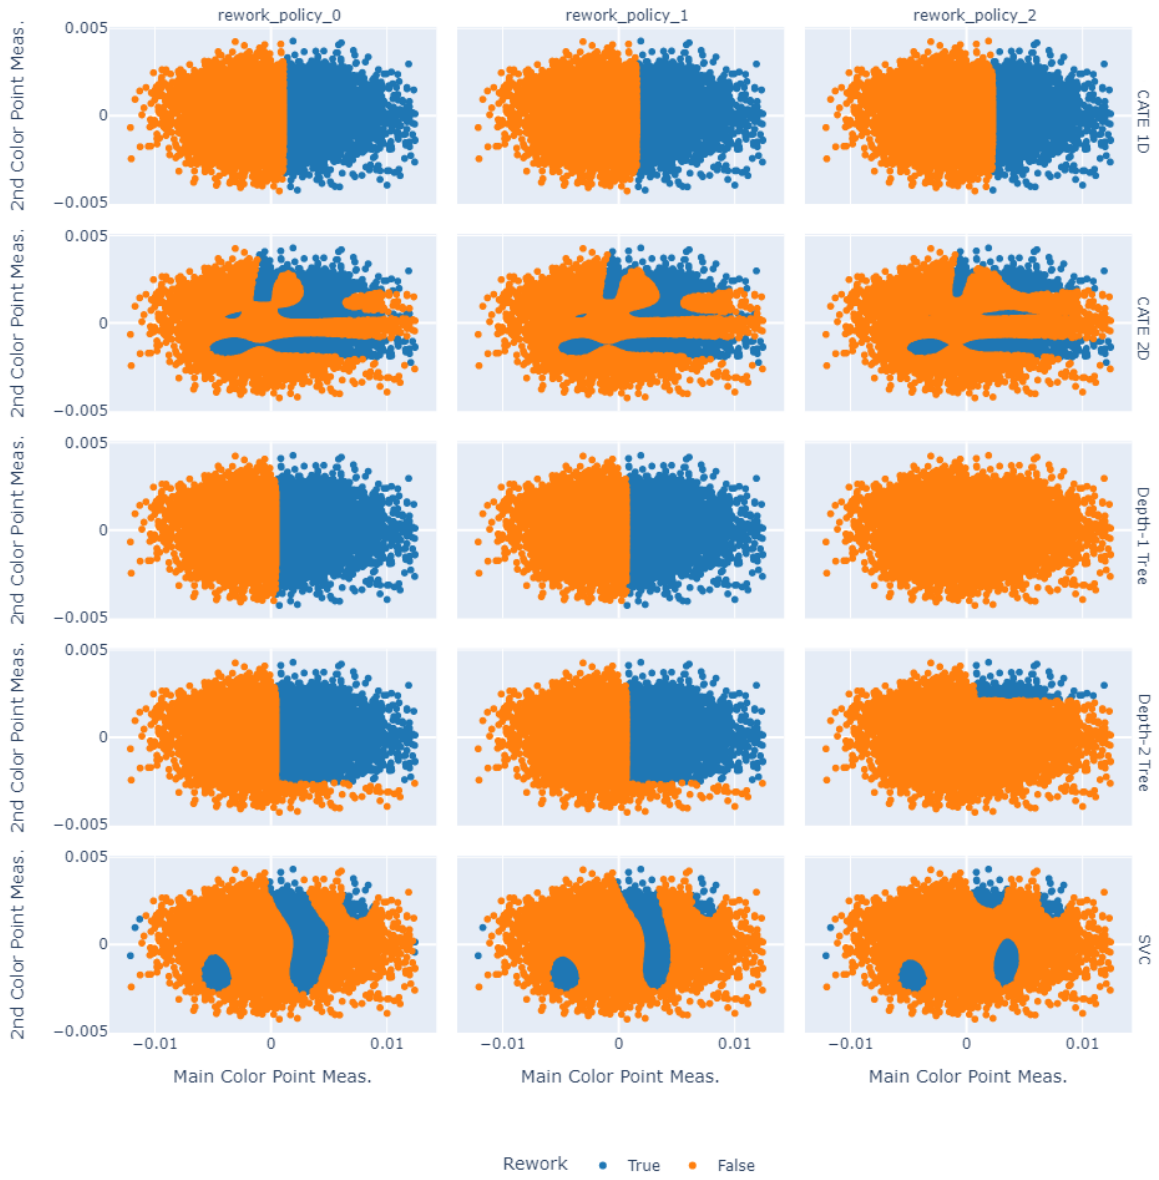


Figure 13: Comparison of the estimated policies with the nuisance components estimated by linear models.

A.3 Policies Based on Confidence Intervals

As additional results, we derive policies based on the conditional average treatment effect lower confidence interval. We construct a pointwise-valid confidence interval via a gaussian multiplier bootstrap as proposed in Semenova and Chernozhukov (2021) at the level 2α . We then evaluate the lower bound to derive a rework cut-off, for which we are certain that the effect is above the desired threshold $\gamma \in \{0.01, 0.03, 0.05\}$ with probability $1 - \alpha$. We choose $\alpha = 5\%$.

Table 5 displays the share of chip lots treated under the conservative policies. We can deduct that compared to Table 1, less chip lots are assigned to be reworked.

Method	Policy 0	Policy 1	Policy 2
CATE 1D	0.3446	0.2754	0.2067
CATE 2D	0.2152	0.0852	0.0459

Table 5: Share of chip lots treated under the conservative policies.

Unsurprisingly, this leads to an increase of GATE (Table 6) in the group of treated, compared to Table 2.

Method	Policy 0	Policy 1	Policy 2
CATE 1D	0.0758	0.0862	0.0967
CATE 2D	0.0797	0.0950	0.1328

Table 6: Grouped average treatment effect for the group of chip lots treated under the conservative policies.

In the case of the one-dimensional CATE, we believe this approach to be a possibility to use an even more conservative cut-off and still increase the overall yield. Concerning the two-dimensional CATE, we are careful in interpretation, as using the confidence interval might overfit to a large extent on small areas with few observations but a potentially high value. The values of these policies are displayed in Table 7 (in comparison to Table 3).

Method	0	1	2
CATE 1D	0.0261	0.0237	0.0200
CATE 2D	0.0171	0.0081	0.0061

Table 7: Value of the conservative CATE policies.

Additionally, we illustrate in Figure 14 the influence of the little overlap outside the interval $C_m \in [-0.005, 0.005]$. While the CATE is approximately constant, the intervals are becoming increasingly large, making the analysis more difficult due to statistical uncertainty.




Figure 14: CATE over the main color measurement for the full space of C_m values.

Appendix B. Results for Product Type “P2”

In this Section, we will present the analysis for product type “P2”. Overall, the results are very similar to the main analysis (“P1”), as the process is the same but only the product properties differ to some extent. Also for this product we would suggest a cut-off along the main component.

B.1 Results

Method	Policy 0	Policy 1	Policy 2
CATE 1D	0.4848	0.4087	0.3455
CATE 2D	0.5106	0.3856	0.3127
Depth-1 Tree	0.4923	0.4907	0.3708
Depth-2 Tree	0.5624	0.4794	0.3013

Table 8: Share of chip lots treated under the estimated policies for product type “P2” (Share in the observations: 0.2472).




Figure 15: Comparison of the estimated policies for product type “P2”.

Method	Policy 0	Policy 1	Policy 2
CATE 1D	0.0818	0.0880	0.0961
CATE 2D	0.0777	0.0950	0.1112
Depth-1 Tree	0.0819	0.0821	0.0951
Depth-2 Tree	0.0772	0.0865	0.1128

Table 9: GATEs for the estimated policies for product type “P2”. (GATE in the group of observed reworked: 0.1040).

Method	Policy 0	Policy 1	Policy 2
CATE 1D	0.0397	0.0360	0.0332
CATE 2D	0.0397	0.0366	0.0348
Depth-1 Tree	0.0403	0.0403	0.0353
Depth-2 Tree	0.0434	0.0415	0.0340

Table 10: Value of the estimated policies for product type “P2” (Value in observed policy: 0.0036).

References

- Susan Athey and Stefan Wager. Policy learning with observational data. *Econometrica*, 89(1):133–161, 2021. doi: <https://doi.org/10.3982/ECTA15732>. URL <https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA15732>.
- Philipp Bach, Victor Chernozhukov, Malte S Kurz, and Martin Spindler. Doubleml—an object-oriented implementation of double machine learning in r. *arXiv preprint arXiv:2103.09603*, 2021.
- Philipp Bach, Victor Chernozhukov, Malte S Kurz, and Martin Spindler. Doubleml—an object-oriented implementation of double machine learning in python. *J. Mach. Learn. Res.*, 23:53–1, 2022.
- Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1):C1–C68, 01 2018. ISSN 1368-4221. doi: [10.1111/ectj.12097](https://doi.org/10.1111/ectj.12097). URL <https://doi.org/10.1111/ectj.12097>.
- Jaehee Cho, Jun Hyuk Park, Jong Kyu Kim, and E. Fred Schubert. White light-emitting diodes: History, progress, and future. *Laser & Photonics Reviews*, 11, 2017.
- S.D.P. Flapper, J.C. Fransoo, R.A.C.M. Broekmeulen, and K. Inderfurth. Planning and control of rework in the process industries: a review. *Production Planning & Control*, 13(1):26–34, 2002. ISSN 0953-7287. doi: [10.1080/09537280110061548](https://doi.org/10.1080/09537280110061548).
- Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. Lightgbm: A highly efficient gradient boosting decision tree. *Advances in neural information processing systems*, 30:3146–3154, 2017.
- Ning Liu, Youngjoo Kim, and Hark Hwang. An optimal operating policy for the production system with rework. *Computers & Industrial Engineering*, 56(3):874–887, 2009. ISSN 0360-8352. doi: <https://doi.org/10.1016/j.cie.2008.09.013>. URL <https://www.sciencedirect.com/science/article/pii/S0360835208002246>. Intelligent Manufacturing and Logistics.
- Judea Pearl. Causal diagrams for empirical research. *Biometrika*, 82(4):669–688, 1995. ISSN 00063444. URL <http://www.jstor.org/stable/2337329>.

F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.

James M Robins and Andrea Rotnitzky. Semiparametric efficiency in multivariate regression models with missing data. *Journal of the American Statistical Association*, 90(429):122–129, 1995.

Donald B. Rubin. Causal inference using potential outcomes: Design, modeling, decisions. *Journal of the American Statistical Association*, 100(469):322–331, 2005. ISSN 01621459. URL <http://www.jstor.org/stable/27590541>.

Vira Semenova and Victor Chernozhukov. Debiased machine learning of conditional average treatment effects and other causal functions. *The Econometrics Journal*, 24(2):264–289, 2021.

Erik Sverdrup, Ayush Kanodia, Zhengyuan Zhou, Susan Athey, and Stefan Wager. policytree: Policy learning via doubly robust empirical welfare maximization over trees. *Journal of Open Source Software*, 5(50):2232, 2020. doi: 10.21105/joss.02232. URL <https://doi.org/10.21105/joss.02232>.

Chi Wang and Qingyun Wu. FLO: fast and lightweight hyperparameter optimization for automl. *CoRR*, abs/1911.04706, 2019. URL <http://arxiv.org/abs/1911.04706>.