

Winner’s Curse Drives False Promises in Data-Driven Decisions: A Case Study in Refugee Matching

Hamsa Bastani

Wharton School, hamsab@wharton.upenn.edu

Osbert Bastani

University of Pennsylvania, obastani@seas.upenn.edu

Bryce McLaughlin

Wharton School, brycemcl@wharton.upenn.edu

A major challenge in data-driven decision-making is accurate policy evaluation—i.e., guaranteeing that a learned decision-making policy achieves the promised benefits. A popular strategy is *model-based* policy evaluation, which estimates a model from data to infer counterfactual outcomes. This strategy is known to produce unwarrantedly optimistic estimates of the true benefit due to the *winner’s curse*. We searched the recent literature on data-driven decision-making, identifying a sample of 55 papers published in the *Management Science* in the past decade; all but two relied on this flawed methodology. Several common justifications are provided: (1) the estimated models are accurate, stable, and well-calibrated, (2) the historical data uses random treatment assignment, (3) the model family is well-specified, and (4) the evaluation methodology uses sample splitting. Unfortunately, we show that no combination of these justifications avoids the winner’s curse. First, we provide a theoretical analysis demonstrating that the winner’s curse can cause large, spurious reported benefits even when all these justifications hold. Second, we perform a simulation study based on the recent and consequential data-driven refugee matching problem. We construct a synthetic refugee matching environment (calibrated to closely match the real setting) but designed so that no assignment policy can improve expected employment compared to random assignment. Model-based methods report large, stable gains of around 60% even when the true effect is zero; these gains are on par with improvements of 22–75% reported in the literature. Our results provide strong evidence against model-based evaluation.

Key words: winner’s curse, data-driven decision-making, policy optimization and evaluation

1. Introduction

The past two decades have seen tremendous gains in the accuracy of statistical and predictive models, due to a confluence of data availability, computational resources, and algorithmic advances. A major application of such estimation is for data-driven decision-making—e.g., targeting interventions to patients, discounts to customers, or even matching refugees to cities—based on observed covariates. In practice, a decision-maker gathers a historical dataset, estimates a model of both outcomes and treatment effects, and then uses this model to learn a good decision-making policy.

Naturally, one would want to know if this learned data-driven policy improves the decision-maker’s objective (e.g., revenue or social welfare) over the existing policy. This requires *policy*

evaluation, which estimates the purported improvement of the data-driven policy over the existing policy. A variety of methods have been proposed for estimating policy improvement with varying assumptions and statistical guarantees. We posit that a key goal of policy evaluation is to guarantee that the true policy improvement lies within a reported confidence interval with high probability—i.e., a decision-maker can have confidence that the proposed policy has benefits consistent with what was promised. We say a policy evaluation methodology is *valid* if it provides this guarantee.

While this sounds straightforward, the key challenge is that we do not observe counterfactual outcomes for decisions that differ from the ones in the historical data. As a consequence, we cannot directly estimate the efficacy of a decision-making policy on a held-out test set the way we can estimate the predictive accuracy of a prediction model.

When the historical dataset uses random treatment assignment (e.g., the data was collected via a randomized controlled trial (RCT)), then the gold standard policy evaluation methodology is inverse probability weighting (IPW) (Horvitz and Thompson 1952, Rosenbaum and Rubin 1983) on a held-out test set. Intuitively, IPW only evaluates decisions when the chosen decisions align with those in the historical data, altogether avoiding the need to estimate counterfactual outcomes. We call this a *model-free* approach. These estimates can be combined with standard statistical methodologies to establish valid confidence intervals for policy improvement. However, there are several reasons why IPW may not be feasible. When the action space is large or the covariates are high-dimensional, IPW can have very high variance; yet, these are often the settings where data-driven decisions are most desirable. Alternatively, the treatment assignment may not be random, in which case IPW is invalid.¹

In these cases, practitioners often adopt a popular alternative where they use a model to predict counterfactual outcomes, and then use these counterfactual outcomes to estimate performance improvement. We call this a *model-based* approach. The model can range from a simple linear model to a dynamic Markov decision process to a modern machine learning model such as a random forests. The key feature is that they estimate policy improvement by using their estimated model to impute unobserved outcomes under counterfactual treatment decisions.

It is well known that using the same estimated counterfactual model for both policy optimization and policy evaluation results in an issue known as the *winner's curse* (Harrison and March 1984, Smith and Winkler 2006, Andrews et al. 2024, Zrnic and Fithian 2024), where the estimated policy improvements are optimistically biased. Intuitively, the optimization procedure exploits estimation errors, choosing decisions whose predicted quality exceeds their true quality. If we use the same

¹We are also assuming the treatment probabilities are known; they can alternatively be estimated from data, but this strategy relies on the assumption that there are no unobserved confounders.

model to evaluate these decisions, these estimation errors go uncorrected, often predicting substantial performance improvements where none exist. Existing model-based evaluation methodologies do not correct for this unwarranted optimistic bias, making them invalid.

While the winner’s curse is well-known, existing work has only shown it to hold under specific conditions. As a consequence, many papers provide seemingly plausible justifications for using model-based evaluation. To understand the nature of these justifications, we performed a survey of recent *Management Science* papers on data-driven decision-making (detailed in Section 2). Of the 55 papers that report on data-driven policy improvement, all but two use the model-based (rather than model-free) method. Surprisingly, a large number of papers provided no justification. When justifications were provided, they typically fell into one of the following categories:

- **J1: Accurate, stable, and calibrated:** The estimated model achieves high accuracy on a held-out test set, is stable across data subsamples (e.g., bootstrap), and is well-calibrated.
- **J2: Random treatment assignment:** The historical data has known random assignment.
- **J3: Well-specified:** The model family is sufficiently rich to include the true data-generating process (e.g., random forests).
- **J4: Sample splitting:** The evaluation splits the historical data into two disjoint samples and trains a different model on each sample, one for policy learning and one for policy evaluation.

To see why these justifications are plausible, consider the following argument: (1) if either J1 and J2 hold or J3 holds, then the estimated counterfactual outcomes are unbiased or have small bias; (2) if, in addition, J4 holds, then the estimated policy improvement is independent from the learned policy; (3) therefore, the estimated policy improvement is accurate.

Unfortunately, we show both theoretically and through a simulation study that this argument is incorrect. First, we provide two stylized examples showing that the winner’s curse can be arbitrarily large while simultaneously satisfying these justifications. Second, we perform a realistic simulation experiment in the context of the refugee matching problem (Bansak et al. 2018, Ahani et al. 2021) to show that model-based evaluation can report very large spurious effects even when the reported justifications are met. Our results show that the winner’s curse is substantially more pervasive than previously understood. They provide strong evidence against the validity of model-based policy evaluation methods, suggesting that they should not be used.

We highlight the critical need for better policy evaluation methodologies that are valid while addressing the shortcomings of IPW. For instance, recent work has studied promising directions that integrate policy learning and evaluation to provide lower-variance policy improvement estimates while preserving validity—e.g., targeting statistically significant policy improvement instead of purely maximizing expected outcomes (Bastani et al. 2025, Chernozhukov et al. 2025), reducing the dimensionality of the policy class to mitigate the winner’s curse (Banerjee et al. 2025), or rigorously integrating auxiliary datasets with real-world outcome data (Mandyam et al. 2025).

1.1. Our Contributions

Theoretical analysis. Intuitively, the winner’s curse does not arise if the estimated counterfactuals are unbiased and sample splitting is used. However, unbiased estimates are a very strong assumption. In practice, one of two possible shortcomings can arise:

- **Misspecification:** A simple parametric model (e.g., linear) is unlikely to match the true data-generating process, even if it achieves a highly accurate fit on the training set distribution. In this case, even with an infinitely large, high-quality dataset, the estimated model is biased. We provide a simple example showing that misspecification can lead to arbitrarily large, unsubstantiated estimated policy improvements even when justifications J1, J2, and J4 hold. This analysis is detailed in Section 3.
- **Regularized:** One way to avoid misspecification is to use modern non-parametric algorithms such as random forests or gradient boosted machines (GBMs). However, these models must be regularized to avoid overfitting. We show that even when the model family is well-specified, the bias induced by regularization in finite sample can cause the winner’s curse. While this bias is small and stable on the historical data distribution from which the training data is sampled, it can become significant on the shifted data distribution induced by a different policy. We provide an explicit construction via a toy ridge regression model showing it can lead to arbitrarily large, unsubstantiated estimated policy improvements even when all the above justifications (J1, J2, J3, and J4) hold. This analysis is detailed in Section 4.

The literature specifically on the winner’s curse typically focuses on the well-specified setting without sample splitting (Harrison and March 1984, Smith and Winkler 2006); furthermore, they generally rely on the estimated model being noisy (i.e., achieves low accuracy on a held-out test set) and unstable (i.e., parameter estimates can change drastically when the model is trained on new or bootstrapped samples). Our analysis shows that due to regularization, the winner’s curse can occur under much more general conditions. Separately, there has been work showing that misspecification can result in bias (Jiang and Li 2016, Kallus and Uehara 2019); however, they do not show that this bias can be systematically positive, or that it can arise even when the estimated model is highly accurate and stable.

Simulation study. Empirically, we perform a simulation experiment in the context of data-driven refugee matching (Bansak et al. 2018, Ahani et al. 2021), where the goal is to match refugees to capacity-constrained cities to optimize overall employment outcomes. While model-based evaluations are widely used in the literature, we focus on refugee matching both due to widespread recent interest in this problem as well as the fact that it has been deployed across several countries to make placement decisions for thousands of refugees. Both Bansak et al. (2018) and Ahani et al. (2021) use the model-based method to evaluate proposed data-driven refugee matching algorithms,

estimating 22–75% gains in refugee employment outcomes from data-driven placement decisions. We design a simulated dataset that closely matches the setting described in Bansak et al. (2018), but constructed so that no policy can outperform random assignment. Thus, any estimated performance gains are spurious by design. Applying the algorithm and model-based evaluation method from Bansak et al. (2018) to this dataset, we obtain (spurious) estimated gains in employment outcomes of around 60%, which is on par with the estimates reported in their paper. We additionally consider a more sophisticated variation of the model-based evaluation method using bootstrapping (Ahani et al. 2021), and find that it also continues to produce substantial, stable estimated improvements even when the ground truth is null.²

1.2. Related Literature

Policy evaluation methods are broadly categorized into model-free and model-based approaches. Model-free evaluation, rooted in the foundational work of Horvitz and Thompson (1952) and Rosenbaum and Rubin (1983), provides asymptotically unbiased estimates of policy performance. However, these methods exhibit substantial variance, particularly in settings with large action spaces (Saito and Joachims 2022). The granularity required for data-driven decision-making undermines standard policy evaluation mechanisms: as the support of the historical data becomes sparse relative to the covariate/action space, inverse probability weighting (IPW) yields extreme weights and unstable estimates. While doubly robust estimation (notably the augmented IPW estimator, see Robins et al. 1994, Robins and Rotnitzky 1995) improves efficiency, the variance remains problematic in finite samples.

Given these limitations, many studies instead rely on model-based evaluation, which is subject to the winner’s curse. The winner’s curse was originally studied in auction theory, where the winner typically overestimates an item’s value (Capen et al. 1971); the concept was formalized for decision analysis by Harrison and March (1984) and Smith and Winkler (2006). In this work, we rebut common justifications used in the literature to dismiss the winner’s curse.

Recent literature addresses these challenges through three primary mechanisms. First, inference methods account for selection effects, either by constructing confidence intervals conditional on the selection event (Andrews et al. 2024) or by applying corrections (Zrnic and Fithian 2024). Similarly, Gupta et al. (2024) and Xu et al. (2025) propose debiasing in-sample performance using gradient-based and bootstrap-based corrections respectively. Second, variance reduction can be achieved by pooling treatments (Banerjee et al. 2025) or leveraging auxiliary data (Mandyam et al.

² To the best of our knowledge, all published estimates for this problem use some variation of the model-based method. Prior to posting this paper, we shared our critique of model-based evaluation methods in private correspondence with Bansak et al. (2018) and Ahani et al. (2021); in response, Bansak et al. (2018) have started developing model-free estimates to rigorously establish performance gains, and have shared preliminary results that appear promising.

2025). Third, the optimization objective itself can be modified; for example, Swaminathan and Joachims (2015) penalize variance to constrain the proposed policy to be close to the support of the existing policy, so that it can be evaluated downstream. Similar principles underpin “pessimistic learning” in offline reinforcement learning (Kumar et al. 2020). Most recently, Chernozhukov et al. (2025) and Bastani et al. (2025) propose learning policies that target statistical significance on a held-out test set, with the latter explicitly characterizing the Pareto frontier of policies that trade off expected performance and statistical significance.

2. A Survey of the Literature

To understand the extent to which model-based policy optimization is currently used in the literature—and consequently, the prevalence of the winner’s curse—we surveyed recent papers published in *Management Science*. We focused on this journal as a leading outlet for quantitative research at the intersection of data analysis (estimation) and decision-making (optimization).

We constructed a sample of recent papers using the EBSCO Business Complete database. We searched for papers published in the last ten years (2015–2024) that contained keywords related to both data-driven estimation and policy optimization. Specifically, we used the query (data OR estim*) AND (optim* OR policy) and restricted the search to peer-reviewed articles in *Management Science*. This initial search yielded 875 results. To efficiently screen this large corpus, we employed a two-stage classification procedure involving a Large Language Model (LLM) followed by manual verification. In the first stage, we fed the titles and abstracts of the 875 papers to the GPT-5.2 API with high reasoning. We prompted the model to identify papers that likely met three specific criteria: (1) one of the primary contributions involves estimating a model (e.g., demand, welfare) from data; (2) the estimated model is used to optimize a decision (e.g., pricing, routing, allocation); and (3) the abstract explicitly states an estimated improvement resulting from this optimized decision, suggesting that this performance gain is a headline contribution. To calibrate the model, we provided 39 manually labeled examples (both positive and negative) and instructed it to follow the provided reasoning. This automated screening process identified 111 papers as potentially fitting our criteria (classified as “yes” or “unclear”).

In the second stage, we manually downloaded and reviewed the full text of these 111 papers. We retained only those fitting the “estimate-then-optimize” framework where evaluation was performed on historical data with unknown counterfactuals. We further excluded papers where the optimization was purely theoretical without empirical estimation, or where the estimation was not central to the policy construction (e.g., purely descriptive regression analysis followed by a simulation-based analysis).

As summarized in Table 1, this process yielded 55 papers. Of these, 53 (96%) relied on model-based evaluation—specifically, they evaluated the performance of their proposed policy using the

Table 1 Survey of Estimate-Then-Optimize Papers in Management Science

Search Criteria	Count
(1) Initial Search Results (Keyword match, last 10 years)	875
(2) Potentially Relevant (AI Screened via GPT-5.2)	111
(3) Confirmed Estimate-Then-Optimize (Manual Review)	55
(4) Susceptible to Winner’s Curse (Model-Based Evaluation)	53

Notes: Rows 1 and 2 represent a sample of papers identified only via their titles and abstracts. Row 3 represents papers confirmed via manual review of the full text to follow an estimate-then-optimize pipeline. Row 4 indicates the subset of Row 3 that rely on model-based evaluation.

Table 2 Survey of model-based counterfactual practices in estimate-then-optimize papers

	Count (%)
Model-based evaluation	
Optim & eval use same model structure	53 (100%)
Optim & eval use exact same model	35 (66%)
Justifications for model-based counterfactuals	
Accuracy	25 (47%)
Stability	19 (36%)
Calibration	7 (13%)
Sample splitting	4 (8%)
Confidence Intervals for estimated gains	
Confidence interval reported	9 (17%)
Model class used for estimation	
Simple parametric statistical models	35 (66%)
Structural models	13 (25%)
Machine learning models	5 (9%)
Problem domain	
Revenue management	24 (45%)
Healthcare	16 (30%)
Service	11 (21%)
Finance/marketing	2 (4%)

Notes: “Simple parametric” includes linear/logit/MNL (both offline and online) as well as related classical econometric specifications; “structural” includes MDP/queueing/SIR; “machine learning” includes KNN/KDE/DQN and uncertainty sets.

same data generating process or specific dataset used for estimation. As demonstrated in the rest of this paper, this methodology creates a direct channel for the winner’s curse to inflate performance estimates. This survey suggests that the issue is not merely a theoretical curiosity but a pervasive feature of the current state of data-driven decision-making in the field.

Table 2 summarizes the specific characteristics of these 53 papers. We start with the mechanics of the evaluation. Thirty-five of the 53 papers (66%) evaluate the performance of their optimized policy using the *exact same* estimated model that was used to generate the policy. This is the most

severe form of the winner’s curse, as the optimizer is free to exploit every residual error in the model. The remaining papers generally use the same model structure but re-estimate parameters, which, as we show in Section 4, does not eliminate the bias.

How do these papers justify this approach? Most do not. Among those that do, the most commonly cited defense is predictive accuracy. Twenty-five papers (47%) argue that because the underlying estimated model achieves high accuracy (e.g., low MSE on a held-out test set), the resulting policy evaluation is valid. Nineteen papers (36%) appeal to stability (e.g., the estimated gains are stable across bootstrapped subsamples). Calibration of predictions is cited less frequently, appearing in only seven papers (13%). Sample splitting—which ensures that the evaluation and optimization procedures do not re-use the same data—is rare, appearing in only four papers (8%). As we demonstrate in Sections 3 & 4, none of these justifications—accuracy, stability, calibration, or sample splitting—are sufficient to rule out spurious gains from the winner’s curse.

The vast majority of papers do not address the uncertainty of the estimated gains at all. Only 9 papers (17%) report confidence intervals for their policy improvements. None attempt to construct these intervals in a way that accounts for the optimization process itself. This lack of statistical inference is particularly troubling given the magnitude of the claimed improvements, which may rely on exploiting noisy predictions from tail events or rare subpopulations.

Two additional points are worth noting regarding the model classes and problem domains. First, the winner’s curse is not limited to “black box” machine learning methods. Thirty-five of the papers (66%) use simple parametric models (such as linear or logit regressions, either via offline or online learning), 13 (25%) use structural models (such as MDPs or queueing), and only five (9%) use blackbox machine learning models (such as KNNs or DQNs). This suggests that the issue is fundamental to the estimate-then-optimize paradigm rather than the adoption of blackbox machine learning methods. Second, the practice is widespread across domains. Revenue management (45%) and healthcare (30%) are the most represented fields, likely due to the natural fit of optimization techniques in these areas. In short, our survey suggests that model-based evaluation is the default standard in the literature, despite its inherent susceptibility to the winner’s curse.

3. Misspecification Causes the Winner’s Curse

In this section, we provide a theoretical construction where misspecification drives the winner’s curse even when justifications J1 (accurate, stable, and calibrated model), J2 (historical data uses random treatment assignment), and J4 (sample splitting) all hold. While our example is stylized, as we discuss in Section 3.3, the key aspects of our example that drive the winner’s curse—namely, the exploitation of estimation errors in regions of sparse data—correspond to natural behaviors of simple estimators (e.g., ordinary linear regression) in high-dimensional settings.

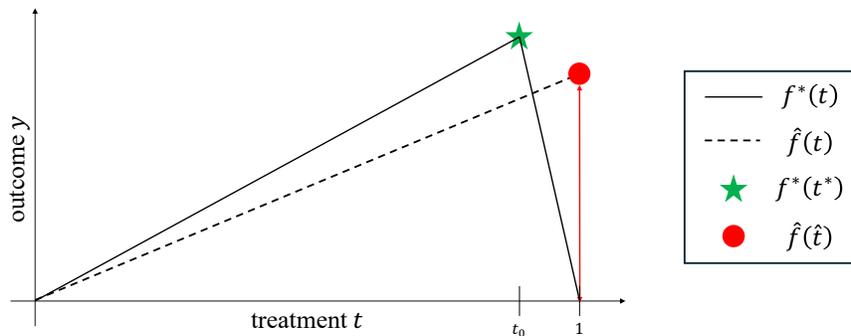


Figure 1 Illustration of stylized example. The x -axis is the treatment t , the y -axis is the performance outcome y , the solid black line is $f^*(t)$, the dashed black line is $\hat{f}(t)$, the green star denotes $f^*(t^*)$ at $t^* = \arg \max_{t \in [0,1]} f^*(t)$, the red circle denotes $\hat{f}(\hat{t})$ at $\hat{t} = \arg \max_{t \in [0,1]} \hat{f}(t)$, and the red line denotes the optimistic bias $\hat{f}(\hat{t}) - f^*(\hat{t})$ due to the winner's curse.

3.1. Theoretical Analysis

We consider a simple setting with a single real-valued treatment $t \in [0, 1]$ and an outcome $y \in \mathbb{R}$, without covariates. We consider a noiseless true model $y = f^*(t)$, where f^* is a piecewise linear function with two hyperparameters $t_0 \in (0, 1)$ and $y_{\max} \in \mathbb{R}$:

$$f^*(t) = \begin{cases} \frac{t y_{\max}}{t_0} & \text{if } t \leq t_0 \\ \frac{(1-t)y_{\max}}{1-t_0} & \text{if } t > t_0. \end{cases}$$

This function is illustrated as the solid black line in Figure 1.

To estimate this model, we consider historical data collected using a uniform distribution over treatments $t \sim \text{Uniform}([0, 1])$. Specifically, a single example in our historical dataset is a pair (t, y) , where $y = f^*(t)$. For simplicity, we consider the infinite data limit, where we observe the exact distribution over t . We assume we are fitting a linear function βt , a standard model family that achieves high out-of-sample accuracy in this setting.

LEMMA 1. *The ordinary least squares (OLS) estimate of $\beta \in \mathbb{R}$ is*

$$\hat{\beta} = \frac{(1 + t_0)y_{\max}}{2}.$$

Proof of Lemma 1 Note that the OLS estimator is $\hat{\beta} = \mathbb{E}[t^2]^{-1} \mathbb{E}[t f^*(t)]$. We have

$$\mathbb{E}[t^2] = \int_0^1 t^2 dt = \frac{1}{3},$$

and

$$\begin{aligned} \mathbb{E}[t f^*(t)] &= \int_0^{t_0} \frac{t^2 y_{\max}}{t_0} dt + \int_{t_0}^1 \frac{t(1-t)y_{\max}}{1-t_0} dt \\ &= \frac{t_0^2 y_{\max}}{3} + \int_0^{1-t_0} \frac{s(1-s)y_{\max}}{1-t_0} ds \end{aligned}$$

$$\begin{aligned}
&= \frac{t_0^2 y_{\max}}{3} + \frac{(1-t_0)y_{\max}}{2} - \frac{(1-t_0)^2 y_{\max}}{3} \\
&= \frac{(1+t_0)y_{\max}}{6}.
\end{aligned}$$

The claim follows. \square

Thus, our model is $\hat{f}(t) = \hat{\beta}t$; this function is visualized as the dashed line in Figure 1. Our first result characterizes the winner’s curse for this model.

PROPOSITION 1. *The optimistic bias from the winner’s curse is $\frac{(1+t_0)y_{\max}}{2}$.*

Proof of Proposition 1 The best treatment according to our model is $\hat{t} = \arg \max \hat{f}(t) = 1$. It has estimated outcome $\hat{y} = \hat{f}(\hat{t}) = \frac{(1+t_0)y_{\max}}{2}$ and true outcome $\tilde{y} = f^*(\hat{y}) = 0$, so the bias is $\hat{y} - \tilde{y} = \frac{(1+t_0)y_{\max}}{2}$. \square

For $t_0 \geq \frac{1}{2}$, we have $\hat{y} - \tilde{y} \geq \frac{y_{\max}}{2}$; thus, the winner’s curse can be arbitrarily large. Note that the estimated performance is comparable to the true optimal performance (the optimal treatment is $t^* = \arg \max_{t \in [0,1]} f^*(t) = t_0$ and its true outcome is $y^* = f^*(t^*) = y_{\max}$). Yet, the realized performance is zero—i.e., the estimated outcome is entirely illusory due to the winner’s curse.

Crucially, this happens even though the model is highly accurate on the historical data distribution (i.e., achieving arbitrarily low MSE).

PROPOSITION 2. *Letting $t_0 = 1 - \epsilon$, if $\epsilon \leq \frac{1}{2}$, then the mean-squared error satisfies*

$$MSE(\hat{f}) \leq 2\epsilon y_{\max}^2 + \frac{\epsilon^2 y_{\max}^2}{3}.$$

Proof of Proposition 2 Note that

$$\begin{aligned}
MSE(\hat{f}) &= \int_0^{1-\epsilon} \left(\frac{(2-\epsilon)ty_{\max}}{2} - \frac{ty_{\max}}{1-\epsilon} \right)^2 dt + \int_{1-\epsilon}^1 \left(\frac{(2-\epsilon)ty_{\max}}{2} - \frac{(1-t)y_{\max}}{\epsilon} \right)^2 dt \\
&= \left(\frac{1}{1-\epsilon} + \frac{1}{2} \right)^2 \cdot \frac{(1-\epsilon)^3 \epsilon^2 y_{\max}^2}{3} + \int_0^\epsilon \left(\frac{(2-\epsilon)(1-s)y_{\max}}{2} - \frac{sy_{\max}}{\epsilon} \right)^2 ds \\
&\leq \frac{\epsilon^2 y_{\max}^2}{3} + 2\epsilon y_{\max}^2,
\end{aligned}$$

as claimed. \square

Intuitively, the discrepancy between high accuracy and large winner’s curse occurs because the optimized treatment induces a shift in the data distribution—the estimated model performs significantly worse on this shifted distribution.

3.2. Contextualizing Our Result

Our model illustrates that even when justifications J1, J2, and J4 in Section 1 hold, the winner’s curse can result in arbitrarily large, spurious policy improvement estimates. Here, we connect our example to these justifications:

- **J1: Accurate, stable, and calibrated:** Our model \hat{f} is perfectly stable since we are in the limit of infinite data. Furthermore, as $t_0 \rightarrow 1$ (equivalently, $\epsilon \rightarrow 0$), it is highly accurate. A highly accurate model is by construction well-calibrated, since it can simply always express confidence in its prediction.
- **J2: Random treatment assignment:** Our historical data is generated using random treatment assignment (i.e., uniformly random).
- **J4: Sample splitting:** Because we take the limit of infinite data, we can retrain the model under independent data splits and the estimated models on each split will be identical. Thus, sample splitting does not affect the policy improvement estimate.

3.3. Connection to Practice

While our example is highly stylized, it proves that the justifications described in Section 1 (except the well-specified justification, J3) are insufficient to rule out the winner’s curse from a theoretical perspective. To connect this stylized example to practical failure modes, consider the two key factors that drive the winner’s curse:

1. Because the model family is misspecified, the estimated model produces biased counterfactual estimates even in the limit of infinite data.
2. The bias occurs in a small region of the input space (i.e., $t \in (t_0, 1]$), allowing the model to achieve very low MSE on the observed data distribution.

Then, the optimizer can exploit the high-error region of the input space to maximize the predicted outcome, leading to the winner’s curse.

One notable feature of our example is that the outcome goes from y_{\max} to zero with a tiny change in treatment, which might be unlikely to happen in practice. However, this construction was only necessary since we were using a simple, one-dimensional treatment and no covariates. In realistic examples, the input space is typically high-dimensional, providing significantly more room for this kind of bias to occur. For instance, in the refugee setting, the goal is to assign refugees to cities to maximize employment rate (e.g., after 90 days). In this problem, each city is a discrete treatment $t \in \mathcal{T} = \{1, \dots, k\}$. This problem additionally involves targeting the treatment based on individual covariates $x \in \mathcal{X} \subseteq \mathbb{R}^d$, such as country of origin, languages spoken, etc. The outcome for a single individual is an indicator for employment $y \in \{0, 1\}$. In this case, a policy has form $\pi : \mathcal{X} \rightarrow \mathcal{T}$. A standard way to learn such a policy called the *plug-in approach* is to first estimate a model $f^\beta : \mathcal{X} \times \mathcal{T} \rightarrow [0, 1]$ mapping covariate-treatment pairs to outcome probabilities. This function induces a policy

$$\pi^\beta(x) = \arg \max_{t \in \mathcal{T}} f^\beta(x, t),$$

i.e., choose the best treatment according to our estimated model.

Consider using a low-capacity model family—e.g., linear models $f^\beta(x, t) = \beta^\top \phi(x, t)$ over a feature map $\phi(x, t) \in \mathbb{R}^m$. If the feature map is too simple, then even for the optimal linear model β^* , some covariate-treatment pairs will naturally be upwards biased—i.e., $f^{\beta^*}(x, t) > \mathbb{E}[y^* | x, t]$, where y^* is the true employment outcome—and some will naturally be downwards biased. Even if this average bias is small (or zero), the optimizer can systematically exploit the upwards-biased covariate-treatment pairs to obtain spurious improvements. Furthermore, as the feature dimension m becomes larger, we might expect the winner’s curse to become larger since there are more opportunities for bias to arise.

4. Regularization Causes the Winner’s Curse

In Section 3, we showed that misspecification can result in the winner’s curse. Thus, one may hope that the winner’s curse is mitigated when using a high-dimensional (e.g., LASSO) or non-parametric model (e.g., random forest); then, we may argue that it is likely that the true data generating process lies in the model family, satisfying J3.

Unfortunately, we show that the winner’s curse can still arise when the estimation problem is well-specified as long as the model is *regularized*. This case is substantially more delicate—indeed, it is true that in the well-specified setting (along with the other justifications), the winner’s curse disappears asymptotically. However, in finite sample, these models must be regularized to avoid arbitrary overfitting, which creates bias in finite sample. Surprisingly, this bias can actually be *stable*—i.e., the estimated model is biased in the same way with high probability over data sub-samples. Intuitively, regularization is typically applied in a systematic way, forcing “high variability” components of the model family to zero until a sufficient amount of data is available to estimate them accurately.

4.1. Theoretical Analysis

We show via a simple well-specified linear model that in fact, regularization bias is sufficient to result in the winner’s curse. Specifically, consider a linear model with a binary treatment $t \in \{0, 1\}$, a single continuous covariate $x \in \mathbb{R}$, and an interaction term tx . We assume that $t \sim \text{Bernoulli}(p)$ (for some $p \in [0, 1]$) and $x \sim \mathcal{N}(0, 1)$ are independent random variables. Also, assume that the noise distribution is $\epsilon \sim \mathcal{N}(0, 1)$. We can interpret this setting in the context of our refugee matching case study by considering two locations A, B , with t indicating whether to assign a refugee to location A (if $t = 0$) or B (if $t = 1$). Each individual is represented by a single covariate x .

To estimate a useful targeting model, we need to have interaction terms between t and x ; otherwise, a linear model will conclude that a single treatment is best for all individuals. Thus, we consider the following feature map (which we assume to be well-specified):

$$\phi(x, t) = [t \ x \ tx]^\top,$$

which includes a single interaction term tx . In this section, we distinguish between *covariates* (i.e., the individual attribute x) and *features* (i.e., the components of the feature map $\phi(x, t)$ encoding covariate-treatment pairs into a real-valued vector). Next, we assume the true parameters are

$$\beta^* = [0 \ b \ 0]^\top.$$

Note that under these parameters, there is no value to targeting—the outcome only depends on x , and not on the interaction tx .

Now, suppose that $p \approx 1$; that is, in our historical data, refugees are overwhelmingly assigned to location B (e.g., due to stringent capacity constraints in A). In this case, there are very few refugees assigned to location A . In the data, we observe that the outcome y is strongly correlated with both x and tx (since x and tx are also strongly correlated), but we cannot distinguish which of these two features is causal; thus, we cannot identify their coefficients. Mathematically, this is because the true covariance matrix

$$\Sigma = \mathbb{E}_{x,t}[\phi(x, t)\phi(x, t)^\top] = \begin{bmatrix} \mathbb{E}_{x,t}[t^2] & \mathbb{E}_{x,t}[tx] & \mathbb{E}_{x,t}[t^2x] \\ \mathbb{E}_{x,t}[tx] & \mathbb{E}_{x,t}[x^2] & \mathbb{E}_{x,t}[tx^2] \\ \mathbb{E}_{x,t}[t^2x] & \mathbb{E}_{x,t}[tx^2] & \mathbb{E}_{x,t}[t^2x^2] \end{bmatrix} = \begin{bmatrix} p & 0 & 0 \\ 0 & 1 & p \\ 0 & p & p \end{bmatrix},$$

has $\lambda_{\min}(\Sigma) \rightarrow 0$ as $p \rightarrow 1$, highlighting the strong correlation between x and tx . As a consequence, any *unbiased* estimate of β^* (i.e., via OLS) will be highly unstable. Specifically, while we can accurately estimate β_1^* from just a few samples, estimating β_2^*, β_3^* requires many more samples due to insufficient variation in the training data. This is not a problem for prediction—we can still identify the sum $\beta_2^* + \beta_3^* = b$. If our goal is only to achieve good out-of-sample accuracy on the distribution of the historical data, then identifying $\beta_2^* + \beta_3^*$ is sufficient, since $t = 1$ with high probability, so

$$\beta_2^*x + \beta_3^*xt = (\beta_2^* + \beta_3^*)x = (\beta_2^* + \beta_3^*)tx.$$

We can obtain a stable estimate of $\beta_2^* + \beta_3^*$ by using regularization, specifically, ridge regression. Intuitively, regularization forces the ridge regression estimator to evenly distribute b between β_2^* and β_3^* ; indeed, consider the parameter vector

$$\bar{\beta}^{(\lambda)} = [0 \ b/2 \ b/2]^\top.$$

We can show that the ridge regression estimator $\hat{\beta}^{(\lambda)}$ concentrates tightly to $\bar{\beta}^{(\lambda)}$, implying that it is stable. Furthermore, $\bar{\beta}^{(\lambda)}$ satisfies the desired relation $\bar{\beta}_2^{(\lambda)} + \bar{\beta}_3^{(\lambda)} = b$; thus, this estimator has high accuracy. Specifically, define the *excess mean-squared error (MSE)* to be the generalization error relative to the true model (typically estimated on a held-out test set):

$$\mathcal{E}(\beta) = \mathbb{E}_{x,t}[(f^\beta(x, t) - f^{\beta^*}(x, t))^2],$$

where $\mathbb{E}_{x,t}$ is the expectation with respect to the training covariates $X \sim \mathcal{N}(0,1)^n$ and training treatments $T \sim \text{Bernoulli}(p)^n$. Then, $\mathcal{E}(\bar{\beta}^{(\lambda)})$ can be small (e.g., on the order of $1/\sqrt{n}$) even when $\|\bar{\beta}^{(\lambda)} - \beta^*\|_2$ is large (e.g., on the order of 1).

However, our goal is policy learning—thus, it is not enough to achieve good accuracy on the observed data distribution. Instead, we need to obtain good accuracy on the eventual data distribution induced by using an alternative policy $\pi : \mathcal{X} \rightarrow \mathcal{T}$ to assign treatments. Unfortunately, this requires that we identify β_2^* and β_3^* separately; the winner’s curse arises from this disconnect. In particular, note that $\pi^{\bar{\beta}^{(\lambda)}}(x) = \mathbb{1}(x \geq 0)$. Because the interaction term tx has estimated coefficient $\bar{\beta}_3^{(\lambda)} = b/2$, the policy believes it is beneficial for units with $x \geq 0$ to receive treatment $t = 1$; conversely, for $x \leq 0$, it believes it is beneficial for units to receive treatment $t = 0$.

DEFINITION 1. Given $\beta, \beta' \in \mathcal{B}$, the *policy improvement* of β when evaluated using β' is

$$\mathcal{P}(\beta; \beta') = \mathbb{E}_x[f^{\beta'}(x, \pi^\beta(x))],$$

and the *optimism bias* is

$$\mathcal{R}(\beta; \beta') = \mathcal{P}(\beta; \beta') - \mathcal{P}(\beta; \beta^*).$$

This optimistic bias for a unit with $x \geq 0$ is $bx/2$, and the overall optimistic bias of $\pi^{\bar{\beta}^{(\lambda)}}$ is

$$\mathcal{R}(\bar{\beta}^{(\lambda)}; \bar{\beta}^{(\lambda)}) \approx \mathbb{E}_x \left[\frac{bx}{2} + \frac{x\mathbb{1}(bx \geq 0)}{2} \right] - \mathbb{E}_x \left[\frac{bx}{2} + \frac{btx}{2} \right] = \frac{b}{2} \cdot \mathbb{E}_x[x\mathbb{1}(x \geq 0)] = \frac{b}{2\sqrt{2\pi}}.$$

Furthermore, note that stability of $\hat{\beta}^{(\lambda)}$ implies stability of the optimism bias (formalized in Proposition EC.3); thus, using sample splitting or a bootstrapped estimate would reliably produce the same optimism bias. We formalize these results in the following proposition.

PROPOSITION 3. Given $\alpha \geq 32$ and $\delta \in \mathbb{R}_{>0}$, let $\lambda = \alpha\sqrt{18\log(8/\delta)/n}$ and $p = 1 - \Delta/2$. Assume that $n^{1/4} \geq 4\sqrt{\log(8/\delta)}/b$. Then, with probability at least $1 - \delta$, all of the following hold:

- **Accuracy:** The MSE of our estimator satisfies

$$\mathcal{E}(\hat{\beta}^{(\lambda)}) \leq 6b^2\alpha\sqrt{\frac{18\log(8/\delta)}{n}} + \frac{72\log(8/\delta)}{n}.$$

- **Stability:** Our estimated model satisfies

$$\|\hat{\beta}^{(\lambda)} - \hat{\beta}^{(\lambda)'}\|_2 \leq \frac{8b}{\alpha} + 4\sqrt{\frac{\log(8/\delta)}{\alpha\sqrt{n}}},$$

where $\hat{\beta}^{(\lambda)}$ and $\hat{\beta}^{(\lambda)'}$ are trained on i.i.d. datasets. Our estimated policy improvement satisfies

$$|\mathcal{P}(\hat{\beta}^{(\lambda)}; \hat{\beta}^{(\lambda)'}) - \mathcal{P}(\hat{\beta}^{(\lambda)''}; \hat{\beta}^{(\lambda)'''})| \leq R \quad \text{where} \quad R = \frac{24b}{\alpha} + 12\sqrt{\frac{\log(8/\delta)}{\alpha\sqrt{n}}} + \frac{128}{\alpha^2} + \frac{32\log(8/\delta)}{\alpha b^2\sqrt{n}},$$

where $\hat{\beta}^{(\lambda)''}$ and $\hat{\beta}^{(\lambda)'''}$ are similarly trained on i.i.d. datasets.

- **Winner’s curse:** *The optimism bias satisfies*

$$\mathcal{R}(\hat{\beta}^{(\lambda)}; \hat{\beta}^{(\lambda)'}) \geq \frac{b}{2\sqrt{2\pi}} - R.$$

We give a proof in Appendix EC.1. If we take $\alpha = n^{1/4}$, then for any choice of b , as n becomes sufficiently large, then the MSE goes to zero and the estimator is perfectly stable, whereas the optimism bias converges to $b/(2\sqrt{2\pi})$. Thus, accuracy and stability cannot definitively rule out the winner’s curse.

4.2. Contextualizing Our Result

Our model illustrates that even when all of the justifications in Section 1 hold, the winner’s curse can result in arbitrarily large, spurious policy improvement estimates. Here, we connect our example to these justifications:

- **J1: Accurate, stable, and calibrated:** Our model $\hat{\beta}^{(\lambda)}$ is accurate (since $\mathcal{E}(\hat{\beta}^{(\lambda)})$ is small for sufficiently large n), stable (since it concentrates for sufficiently large n and taking $\alpha = n^{1/4}$), and well-calibrated (since it can always express valid confidence in its prediction).
- **J2: Random treatment assignment:** Our historical data is generated using random treatment assignment (i.e., randomly assign to location A with probability $1 - p$ and to location B with probability p , where p is known).
- **J3: Well-specified model family:** The model family contains the true model β^* .
- **J4: Sample splitting:** Our result uses sample splitting for policy evaluation.

Yet, by taking b large, we can obtain arbitrarily large unwarranted optimism in our policy improvement estimates.

4.3. Connection to Practice

While our example is stylized, it proves that all of the justifications described in Section 1 are insufficient to rule out the winner’s curse from a theoretical perspective. In this case, the winner’s curse is driven by a very similar mechanism as in Section 3; the key difference is that the bias arises from regularization:

1. Because the model family is regularized, the estimated model produces stable but biased counterfactual estimates.
2. The bias occurs in a low-probability region of the input space (i.e., $t = 0$), allowing the model to achieve very low MSE on the historical data distribution.

As before, the optimizer can exploit the high-error region of the input space to maximize the predicted outcome, leading to the winner’s curse.

In realistic scenarios, low-probability regions of the input space can include not just rare treatments (as in our stylized example), but also individual covariates that rarely occur (e.g., in our

refugee example, a rare country of origin). Our setting is very low-dimensional and used a single, very rare treatment. In practice, the feature dimension m is usually much larger, so it is actually much more likely that there are features $\phi(x, t)_i$ along which variation in the historical data is low, thereby requiring regularization to avoid instability and ensure accuracy. For these features, the optimizer can exploit biases that arise from regularization to achieve spurious improvements.

Beyond linear models, model families such as random forests and gradient boosted machines (GBMs) also perform implicit regularization to avoid overfitting the data. For instance, random forests average over a large number of decision trees, thereby regularizing estimates to zero when uncertain. Alternatively, GBMs rely on fitting low-capacity base models such as shallow decision trees, thereby inheriting the biases of the base model family. Thus, these kinds of models may also exhibit the winner’s curse; our simulation study in Section 5 illustrates this possibility

5. Case Study on Refugee Matching

We now provide empirical evidence that the justifications described in Section 1 do not guarantee accurate policy evaluation using a synthetic environment modeled after the refugee matching problem. Resettlement agencies attempt to settle incoming refugees in locations that maximize their welfare, often measured in the short-run by employment outcomes. To improve their matching of refugees to locations, Bansak et al. (2018) and Ahani et al. (2021) have proposed optimizing future location assignments using employment prediction models trained on historical outcomes. Both papers claim this procedure leads to large gains in employment, using model-based evaluation and justifications J1, J2, and J3. We design a simulation environment calibrated to the U.S.-based setting studied by Bansak et al. (2018), but generate employment outcomes in such a way that no assignment policy can impact the expected employment rate. We train multiple prediction models, use them to estimate optimal policies, then evaluate these policies using both model-based methods (biased by the winner’s curse) and a model-free method (IPW, which is unbiased, but high variance). Despite being built off accurate, calibrated, and stable prediction models (J1), being trained on random historical assignments (J2), and (for some models) being well-specified (J3), all model-based methods falsely report improved employment rates due to the winner’s curse.

5.1. Background

The refugee matching problem provides a structure to analyze and improve assignments of incoming refugees to locations for resettlement. Resettlement agencies assign refugees to locations based on a small set of visible features and a variety of constraints. The authors argue that these constraints introduce random variation to the assignment process that enable cross-location comparisons of refugees with similar features. The goal of refugee matching is to identify feasible assignment policies that improve short-term refugee employment outcomes.

Upon entering the United States, incoming refugees are assigned to a resettlement agency that manages their arrival. Each resettlement agency sets up affiliate networks in various locations around the country to help integrate arriving refugees and provide them various services. Using a limited set of features for each refugee (such as age, nationality, and education), agencies must assign refugees to these locations.

Refugee assignments must satisfy various constraints. Some refugees have prior ties in the United States that dictate their placement. Other refugees require specific services, such as English language education, that only exist in some location's affiliate networks. Moreover, each location has limited resources, creating capacity constraints which may vary over time. These constraints complicate the assignment process, but provide the variation required to estimate the heterogeneous impact of locations on refugees.

Resettlement agencies want to assign refugees in a welfare-maximizing way, subject to their various constraints. In the short-run, few welfare-relevant outcomes are easily trackable, so many agencies record employment outcomes after 90-days as a proxy for a successful refugee-location match. The goal is thus to use past observations of refugee-location matches and their employment outcomes to design an assignment policy that leads to a higher refugee employment rate.

5.2. Optimizing Refugee Matches

Following Bansak et al. (2018) and Ahani et al. (2021), we formulate a simple offline version of the refugee matching problem and describe an estimate-then-optimize method for proposing refugee assignments. Each incoming refugee consists of (known) covariates and (unknown) counterfactual employment outcomes for each location. We use the covariates to estimate employment outcomes for each location, then use an integer program to assign refugees to locations based on these estimates. Additional details are provided in Appendix EC.2.1.

The assignment problem consists of matching N unrestricted refugees to L locations in order to maximize the number of employed refugees. Each refugee $i \in [N]$ is associated with a set of potential outcomes $Y_i(t) \in \{0, 1\} \forall t \in [L]$ which describe whether the refugee would find employment at location t .³ A matching π assigns refugees to locations where $\pi_{it} \in \{0, 1\}$ records whether refugee i is assigned to location t . Each location t can support a maximum of c_t refugees. The optimal matching can then be described by the following integer program.

$$\underset{\pi}{\text{maximize}} \sum_{i=1}^N \sum_{t=1}^L Y_i(t) \pi_{it}$$

³To match the assumption by the original authors, we also assume the standard Stable Unit Value Treatment Assumption which says one refugee's potential outcomes are not impacted by other refugees' location assignments.

$$\begin{aligned}
\text{s.t. } & \sum_{t=1}^L \pi_{it} = 1 \quad \forall i \in [N] \\
& \sum_{i=1}^N \pi_{it} \leq c_t \quad \forall t \in [L] \\
& \pi_{it} \in \{0, 1\} \quad \forall i \in [N], t \in [L].
\end{aligned}$$

However, the counterfactual outcomes under alternative locations for each refugee are unknown. Thus, the authors use a predictive model to estimate these counterfactuals based on refugee covariates X_i (such as age, nationality, and education). Specifically we have access to observations of prior refugee-location matches where for refugee i we observe their covariates X_i , the location they were assigned to $T_i \in [L]$, and their observed employment outcome $Y_i(T_i)$.⁴ These prior observations may have had assignment restrictions, which should be captured by the covariates X_i . We can use these observations to build a prediction model $\hat{\beta}$ where $\hat{\beta}_t(X_i)$ estimates $\mathbb{E}[Y_i(t)|X_i]$ given refugee covariates X_i and a location t . We then propose using the assignment which maximizes the expected number of employments according to the model $\hat{\beta}$, which can be described by the following integer program.

$$\begin{aligned}
& \underset{\pi}{\text{maximize}} \quad \sum_{i=1}^N \sum_{t=1}^L \hat{\beta}_t(X_i) \pi_{it} \\
& \text{s.t. } \quad \sum_{t=1}^L \pi_{it} = 1 \quad \forall i \in [N] \\
& \quad \quad \sum_{i=1}^N \pi_{it} \leq c_t \quad \forall t \in [L] \\
& \quad \quad \pi_{it} \in \{0, 1\} \quad \forall i \in [N], t \in [L].
\end{aligned}$$

We refer to the assignment given by this program as $\pi^{\hat{\beta}}$.

5.3. Evaluating Refugee Matches

Bansak et al. (2018) and Ahani et al. (2021) both evaluate their proposed policies using model-based methods—specifically, they use the *same* model for both policy estimation and policy optimization, which is prone to the winner’s curse as discussed earlier. In this setting, under their assumptions of unconfoundedness and random assignment, a model-free approach (such as IPW), is guaranteed to be valid; however it is prone to high variance due to the size of the action space.

The model-based method of policy evaluation addresses the issue of missing counterfactual outcomes by imputing them. Specifically, it uses the model $\hat{\beta}$ (which optimized the proposed

⁴ We similarly adopt the assumption of unconfoundedness: conditional on the covariates X_i the assignment T_i is independent of the potential outcomes. Later we will assume the exact policy used to generate these matches is known and also assigns positive treatment propensity to each location.

assignment), effectively reporting $\sum_{i=1}^N \sum_{t=1}^L \hat{\beta}_t(X_i) \pi_{it}^{\hat{\beta}}$. Ahani et al. (2021) do the same, but also justify the evaluation using a set of B alternative models $\{\hat{\beta}^{(b)}\}_{b=1}^B$ where model b is built off an independently resampled bootstrap of the original training samples. While straightforward, these model-based approaches may produce biased evaluations if the prediction errors of the model used to propose the assignment $\pi^{\hat{\beta}}$ are correlated with the prediction errors of the model which evaluates the assignment. This is trivially true when $\hat{\beta}$ is reused for evaluation, but may occur even (i) if the evaluation models are built off bootstrapped samples of the training data or (ii) if the evaluation models are built off a completely independent sample. This can occur due to the specification or regularization the prediction models must apply to provide stable insights when training samples are limited.

IPW is a model-free method of policy evaluation which addresses the issue of missing counterfactual outcomes by only evaluating the proposed policy $\pi^{\hat{\beta}}$ when the employment outcome is available. That is we assume each refugee $i \in [N]$ was actually assigned to location T_i resulting in employment outcome $Y_i(T_i)$ where T_i was generated independently from a known distribution where $\Pr(T_i = t) = p_{it} > 0$ for all locations $t \in [L]$. IPW then rescales the outcome data based on how likely each assignment was to occur in the data according to p and estimates the number of employed refugees using the cases the proposed assignment aligns with the observed assignment $\pi_{iT_i}^{\hat{\beta}} = 1$. Thus the expected number of employments under $\pi^{\hat{\beta}}$ according to IPW is

$$\sum_{i: \pi_{iT_i}^{\hat{\beta}}=1} \frac{Y_i(T_i)}{p_{iT_i}}.$$

This estimate of the proposed assignment’s impact is unbiased, but can suffer from excessive variance if p_{iT_i} is small for some refugees. In our environment there are 43 locations, leading to a volatile IPW estimator.

5.4. Synthetic Environment

We develop a synthetic environment to explore the potential impact of the winner’s curse bias on model-based evaluation methods in the refugee matching problem. We tune our setting to mimic the US-based setting of (Bansak et al. 2018) and J2: we generate refugee covariates according to the marginal distributions reported by their supplemental material, then assign the same number of historical locations randomly in their observed proportions. We generate counterfactual employment outcomes in a way that depends on both refugee covariates and locations, but not their interaction, preventing any assignment policy from improving the expected employment rate.

The synthetic environment generates observations of prior refugee-location matches suitable for the policy optimization and evaluation methods described above. That is for each refugee i , our

environment generates covariates X_i , a location T_i , and an employment outcome $Y_i(T_i)$. First, the environment assigns the covariates used by Bansak et al. (2018) (age, gender, education, english-speaking, case restriction, country of origin, arrival year, and arrival month) independently at random according to the marginal distributions reported in their Supplemental Material. Next the location T_i is drawn independently at random from a set of 43 possible locations according to the empirical distribution p that Bansak et al. (2018) reports. We report the covariate and location distributions in Appendix EC.2.2. Finally the employment outcome is drawn independently according to $Y_i(T_i) \sim \text{Ber}(f(X_i, T_i))$ for our causal model f , which dictates the employment probability for a refugee with covariates X_i placed in location T_i .

We generate employment outcomes using a causal model that induces the same expected employment rate for all feasible refugee-location matchings when all locations are at capacity. Specifically we determine the probability a refugee attains employment as

$$f(X_i, T_i) = \frac{1}{2}f_X(X_i) + \frac{1}{2}f_L(T_i),$$

where $f_X \in [0, 1]$ represents the impact of the refugee’s covariates on employment and $f_L \in [0, 1]$ represents the impact of the location on employment. Under a causal model of this form no feasible matching can impact the expected employment rate if all locations are at capacity. We may rearrange the probability of employment across refugees through different assignments, but the lack of covariate-refugee interaction effects prevents us from improving the employment rate via better matchings. The refugee effect, f_X , is determined by first using a logit regression with random coefficients⁵ over refugee covariates and locations to assign an employment outcome to each refugee, then training a random forest regression using the refugee covariates to determine the underlying refugee effect. The location effect, f_L , for each of the 43 locations is drawn independently from a Beta distribution with $\alpha = 1$ and $\beta = 2$.⁶

5.5. Simulation Details

We test the estimate-then-optimize refugee matching method using three prediction model classes, then apply the evaluation methods described in Section 5.3 to the proposed policies. In each case we train the estimation model on 33,000 refugees and evaluate the performance on 1,000 refugees with no prior placement restrictions (as done in Bansak et al. (2018)). We perform each simulation 250 times to obtain histograms of performance estimates for each approach.

⁵ We manually assign the covariate indicating case restriction=free to have a large positive impact. We do this as the test set of Bansak et al. (2018) only includes free case placements and achieves 34% employment while their training set includes both types of cases and achieves 23% employment. To replicate this difference we gave free case placement a positive impact on employment probability.

⁶ We chose these parameters to replicate the employment differences Bansak et al. (2018) observes across locations as faithfully as possible. Employment rates by location are reported in Appendix EC.2.3.

We compare the model-based and model-free methods described above for refugee assignments built using a variety of prediction model classes. First, we examine a LASSO-constrained logit regression over refugee covariates, locations, and their interactions (as done by Ahani et al. (2021)). This model class is very stable and generally performs well out of sample, but is improperly specified based on our causal model f . Second, we examine a family of Honest Random Forests (i.e. causal machine learning (Wager and Athey 2018)); one for each location. The models are known from their ability to produce non-parametric unbiased estimates, but reduce the prediction power of the data in the process. Finally we examine a family of gradient-boosted classification trees (referred to as GBM); one for each location (as done by Bansak et al. (2018)). These models are very powerful, but not very stable and are prone to overfitting.

For each model class our simulation:

1. Generates a training dataset of 33,000 refugees.
2. Trains a prediction model $\hat{\beta}$ of the given model class on the training dataset.
3. Generates a testing data set of 1,000 refugees (all with case restriction = free).
4. Optimizes the assignment of the refugees in the testing dataset according to $\hat{\beta}$, setting location capacities according to the observed number of matches in the testing dataset.
5. Records evaluations of the proposed assignment $\pi^{\hat{\beta}}$ using $\hat{\beta}$ in the model-based method as well as by applying IPW to the testing dataset.
6. Generates 250 bootstrapped versions of the training dataset via resampling, trains a prediction model $\hat{\beta}^{(b)}$ of the given model class for each one, and records evaluation of the proposed assignment $\pi^{\hat{\beta}}$ using $\hat{\beta}^{(b)}$ in the model-based method. Additionally, we simulate 250 IPW evaluations on the same policy by shuffling the location assignments and regenerating employment outcomes. These results are presented in Figure 3.
7. Repeats steps 3. to 5. until 250 estimate-then-optimize assignments have been evaluated. These results are presented in Figure 2.

In addition to histograms which report estimates of the policy’s performance over the 250 runs, we report ROC and calibration curves on a testing dataset for each prediction model in Appendix EC.2.4 to address J1. We generate these curves using SKLearn methods as we do with the logit models and gradient boosted classifiers. We generate Honest Random Forests using the EconML library. We implement the integer optimization problem in CVXPY using the SCIP solver. All simulations are performed in Python.

5.6. Simulation Results

Our simulations reveal that the model-based methods exhibit large biases from the winner’s curse for all three prediction model classes, and that those biases are not eliminated when instead evaluated by models trained on bootstrapped versions of the training dataset. These simulations do

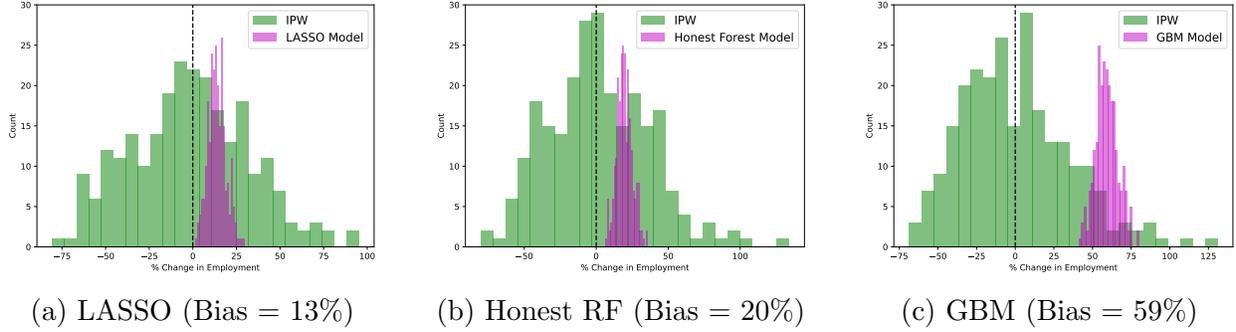


Figure 2 Histograms comparing 250 evaluations according to the prediction model used to optimize the policy with IPW estimates. All prediction models exhibit bias’s due to the winner’s curse while IPW estimation suffers from excessive variance. All evaluations are reported in percent change in employment rate relative to the observed employment rate in the testing dataset.

not prove that the policies learned by Bansak et al. (2018), Ahani et al. (2021) are ineffective; rather, they show that this form of policy evaluation has a very high FPR—consistently promising non-existent gains even when none are present—underscoring the need for valid policy evaluation methodologies in practice.

Figure 2 presents histograms comparing evaluations of the optimized policy’s impact using both the prediction model informing the policy’s construction and IPW estimation. Both LASSO and Honest Random Forests present a moderate winner’s curse bias, while GBM exhibits an extreme bias (likely due to its lack of stability). On the other hand, while IPW estimation remains unbiased regardless of the prediction model used to optimize the policy, its variance, driven by both small sample size and large action space, precludes effective evaluation of the policy’s impact. We caution readers from using these simulations to conclude that LASSO and Honest Random Forest provided less biased evaluation, as alternative parameterizations of our simulation lead to larger biases from LASSO and Honest Random Forest (compared to GBM). Instead, these simulations highlight the need to use provably valid policy optimization and evaluation methodologies that can limit variance in the presence of large action spaces and small sample sizes (see, e.g., Bastani et al. 2025).

Figure 3 presents similar histograms, only now we compare evaluation of the optimized policy using models trained on bootstrapped versions of the training dataset, as performed by Ahani et al. (2021). For LASSO, we find that the bootstrapped models perpetuate as much bias from the winner’s curse as evaluating using the original model used for policy optimization. We believe this follows from misspecification of the stable LASSO model: it is consistently reproducing the predictions of the original model in its bootstraps due to the complicated impact of the refugees’ covariates on employment outcomes. Meanwhile Honest Random Forest and GBM have no obvious misspecification and consequently are able to reduce, but not eliminate, the bias from their evaluations. Instead these techniques reproduce some of the winner’s curse bias that the assignment

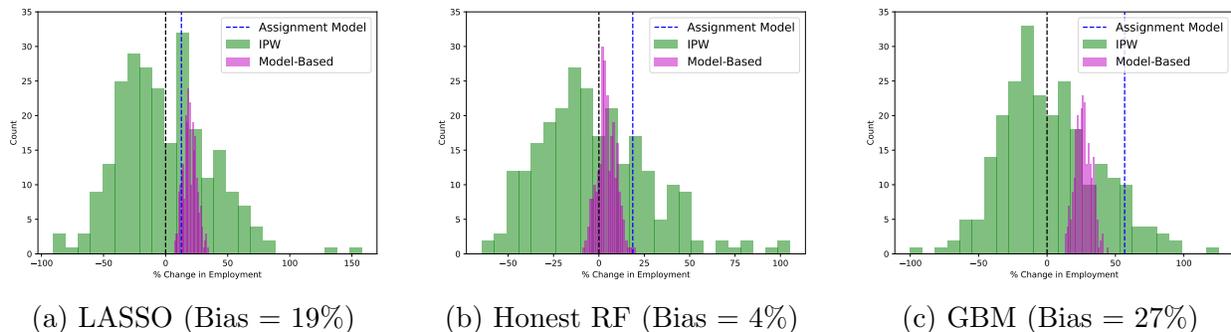


Figure 3 Histograms comparing 250 evaluations according to prediction models built off bootstrapped samples of the training dataset with IPW estimates and direct evaluation by the prediction model which optimized assignments. The stability of the LASSO estimator leads the bootstrapped samples to exhibit a larger winner’s curse than the original prediction model while GBM exhibits a reduced, but still quite large, winner’s curse. Honest Random Forests were able to eliminate most of their winner’s curse bias in this simulation, although in alternate setups we have observed sizable inaccuracies (both positive and negative) indicating the bootstrapping approach cannot generate trustworthy confidence intervals. All evaluations are reported in percent change in employment rate relative to the observed employment rate in the testing dataset.

model suffers due to correlation between the training dataset and its bootstraps. As such we would expect this bias to decrease as the power of the training set increases. However, it is unclear how the size of this bias relates to the size of the training dataset, making it very difficult to use the bootstrapping approach for evaluation.⁷

Overall our simulations highlight that in order to confidently assess complex data-driven assignment policies, we need better policy evaluation methods. Without improved methods, the only way to confidently evaluate these policies is to retain large testing datasets to counteract the variance of IPW estimation, which is often economically infeasible.

6. Discussion and Broader Implications

Our results demonstrate that the winner’s curse is not merely a theoretical curiosity, nor is it confined to obviously misspecified models. Rather, it is a systematic failure mode of model-based policy evaluation in estimate–then–optimize pipelines. In our refugee-matching case study, state-of-the-art model-based methods produced large and stable—but entirely spurious—estimated gains, despite standard checks such as cross-validation and calibration. This is because standard diagnostics validate *average* predictive performance, whereas policy optimization is designed to exploit precisely those parts of the input space where residual errors remain. To clarify why these spurious gains can survive familiar best practices, it is useful to contrast model-based and model-free evaluation across three data regimes.

⁷ We once again caution readers from using these simulations to conclude that Honest Random Forest is the “best” prediction model to use for bootstrapped policy evaluation as other simulation setups using Honest Random Forest models have retained large biases.

The fundamental dilemma in evaluating data-driven policies is a bias–variance trade-off. Model-free methods such as inverse probability weighting (IPW) are asymptotically unbiased but suffer from prohibitively high variance when propensity scores are small or the action space is large. Model-based methods instead impute counterfactual outcomes, reducing variance by smoothing over noise. The problem, however, is that model-based estimates are biased in a systematically optimistic direction. The optimizer preferentially selects actions and subpopulations that appear best under the model—i.e., “wins” are disproportionately driven by favorable estimation error.

In “small data” settings, estimated models are noisy and unstable, discouraging any kind of policy evaluation. At the other extreme, in “large data” regimes with sufficient coverage, model-free estimators become precise, making it possible to accurately evaluate policy improvements.

The most problematic regime—and the one where the winner’s curse is most acute—is “medium data.” Here, the data are sufficient to estimate models that appear accurate, stable, and well-calibrated on average, yet insufficient to support low-variance model-free policy evaluation. In this regime, standard diagnostics create a false sense of security: a practitioner may observe high AUC, stability, and calibration, and conclude the model is a reliable proxy for reality. But policy optimization is effectively an adversarial search process that identifies and exploits residual model errors, especially in regions that are rare under the observed data distribution. Even a model with strong aggregate performance may overestimate outcomes for a small subset of cases; the optimizer will systematically select precisely those cases.

Crucially, because model-free estimates remain high-variance in medium data settings—particularly with large or combinatorial action spaces—they often cannot deliver a decisive rebuttal to a precise (but biased) model-based estimate. The result is a predictable pattern: the optimistic model-based gain looks stable and persuasive, while unbiased checks are too noisy to serve as an effective counterweight. The refugee matching case study exemplifies this regime: roughly 30,000 observations are sufficient to support seemingly reasonable predictive models, yet the sparsity of the assignment space ensures that model-free evaluation is too noisy to robustly validate the optimizer’s purported gains.

Because the medium-data regime is pervasive, the winner’s curse is a systemic threat. Many high-stakes decision problems possess enough data to estimate reasonable statistical models, but not enough to rigorously validate the tails of the distribution where optimized policies concentrate. Thus, we argue against using model-based imputation for policy evaluation, even when the estimated models appear strong by conventional metrics. Effectively doing so requires rigorously integrating model-based data with real-world outcomes (Mandyam et al. 2025).

Rather, emerging research addresses the challenge of the high variance of model-free evaluation by rethinking the learning process itself; for instance, Chernozhukov et al. (2025), Bastani et al. (2025)

propose algorithms that seek statistically guaranteed improvements rather than simple expected value maximization, while Banerjee et al. (2025) demonstrate that constraining the complexity of the policy class can effectively curb the winner’s curse.

Acknowledgments

The authors are grateful to Gad Allon, Jackie Baek, Kirk Bansak, Mohsen Bayati, Ron Berman, Rob Bray, Gerard Cachon, Vishal Gupta, Dean Knox, Elisabeth Paulson, Neha Sharma, Yan-nis Stamotopoulos, Alexander Teytelboym, Stefan Wager, and others for helpful feedback. This research was supported by generous funding from the Wharton AI & Analytics Initiative.

References

- Narges Ahani, Tommy Andersson, Alessandro Martinello, Alexander Teytelboym, and Andrew C Trapp. Placement optimization in refugee resettlement. *Operations Research*, 69(5):1468–1486, 2021.
- Isaiah Andrews, Toru Kitagawa, and Adam McCloskey. Inference on Winners. *The Quarterly Journal of Economics*, 139(1):305–358, January 2024. ISSN 0033-5533, 1531-4650. doi: 10.1093/qje/qjad043. URL <https://academic.oup.com/qje/article/139/1/305/7276491>.
- Abhijit Banerjee, Arun G Chandrasekhar, Suresh Dalpath, Esther Duflo, John Floretta, Matthew O Jackson, Harini Kannan, Francine Loza, Anirudh Sankar, Anna Schrimpf, et al. Selecting the most effective nudge: Evidence from a large-scale experiment on immunization. *Econometrica*, 93(4):1183–1223, 2025.
- Kirk Bansak, Jeremy Ferwerda, Jens Hainmueller, Andrea Dillon, Dominik Hangartner, Duncan Lawrence, and Jeremy Weinstein. Improving refugee integration through data-driven algorithmic assignment. *Science*, 359(6373):325–329, 2018.
- Hamsa Bastani, Osbert Bastani, and Bryce McLaughlin. Beating the winner’s curse via inference-aware policy optimization. *arXiv preprint arXiv:2510.18161*, 2025.
- Edward C Capen, Robert V Clapp, and William M Campbell. Competitive bidding in high-risk situations. *Journal of petroleum technology*, 23(06):641–653, 1971.
- Victor Chernozhukov, Sokbae Lee, Adam M Rosen, and Liyang Sun. Policy learning with confidence. *arXiv preprint arXiv:2502.10653*, 2025.
- Vishal Gupta, Michael Huang, and Paat Rusmevichientong. Debiasing in-sample policy performance for small-data, large-scale optimization. *Operations Research*, 72(2):848–870, 2024.
- J Richard Harrison and James G March. Decision making and postdecision surprises. *Administrative Science Quarterly*, pages 26–42, 1984.
- Daniel G Horvitz and Donovan J Thompson. A generalization of sampling without replacement from a finite universe. *Journal of the American statistical Association*, 47(260):663–685, 1952.
- Nan Jiang and Lihong Li. Doubly robust off-policy value evaluation for reinforcement learning. In *International conference on machine learning*, pages 652–661. PMLR, 2016.

- Nathan Kallus and Masatoshi Uehara. Inherently efficient, stable, and bounded off-policy evaluation for reinforcement learning. *Advances in neural information processing systems*, 32, 2019.
- Aviral Kumar, Aurick Zhou, George Tucker, and Sergey Levine. Conservative q-learning for offline reinforcement learning. *Advances in neural information processing systems*, 33:1179–1191, 2020.
- Aishwarya Mandyam, Jason Meng, Ge Gao, Jiankai Sun, Mac Schwager, Barbara E. Engelhardt, and Emma Brunskill. PERRY: Policy Evaluation with Confidence Intervals using Auxiliary Data, July 2025. URL <http://arxiv.org/abs/2507.20068>.
- 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.
- James M Robins, Andrea Rotnitzky, and Lue Ping Zhao. Estimation of regression coefficients when some regressors are not always observed. *Journal of the American statistical Association*, 89(427):846–866, 1994.
- Paul R. Rosenbaum and Donald B. Rubin. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1):41–55, 1983. ISSN 0006-3444, 1464-3510. doi: 10.1093/biomet/70.1.41. URL <https://academic.oup.com/biomet/article-lookup/doi/10.1093/biomet/70.1.41>.
- Yuta Saito and Thorsten Joachims. Off-policy evaluation for large action spaces via embeddings. *arXiv preprint arXiv:2202.06317*, 2022.
- James E Smith and Robert L Winkler. The optimizer’s curse: Skepticism and postdecision surprise in decision analysis. *Management Science*, 52(3):311–322, 2006.
- Adith Swaminathan and Thorsten Joachims. Batch learning from logged bandit feedback through counterfactual risk minimization. *The Journal of Machine Learning Research*, 16(1):1731–1755, 2015.
- Stefan Wager and Susan Athey. Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113(523):1228–1242, 2018.
- Sikun Xu, Raphael Thomadsen, and Dennis Zhang. The winner’s curse in data-driven decision making: Evidence and solutions. *Available at SSRN 5930537*, 2025.
- Tijana Zrnic and William Fithian. A Flexible Defense Against the Winner’s Curse, 2024. URL <https://arxiv.org/abs/2411.18569>.

Appendix

EC.1. Proofs for Section 4

We begin by proving a number of general results on accuracy and stability of ridge regression; our proof of Proposition 3 relies on the general theory to establish accuracy and stability. Our proof that the winner’s curse happens is specific to the problem construction established in Section 4.

EC.1.1. General Problem Formulation

We begin by formalizing the problem of policy evaluation where the counterfactual outcomes are estimated using a ridge regression model. We consider a covariate space $\mathcal{X} \subseteq \mathbb{R}^d$, a compact space of treatments \mathcal{T} , a parameter space $\beta \in \mathcal{B} \subseteq \mathbb{R}^m$, a feature space $\mathcal{Z} = \mathbb{R}^m$, and a feature map $\phi : \mathcal{X} \times \mathcal{T} \rightarrow \mathcal{Z}$. The *decision objective* is $f^\beta(x, t) = \beta^\top \phi(x, t)$, and the β -*optimal policy* is $\pi^\beta(x) = \arg \max_{t \in \mathcal{T}} f^\beta(x, t)$. We consider a probability measure \mathbb{P}_x (capturing the covariate distribution) and \mathbb{P}_t (capturing the treatment distribution in the historical data), defining the product measure $\mathbb{P}_{x,t} = \mathbb{P}_x \times \mathbb{P}_t$. We assume ϕ is measurable and bounded in expectation: $\mathbb{E}_x [\max_{t \in \mathcal{T}} \|\phi(x, t)\|_2] \leq \phi_{\max}$, and we assume that $\phi(x, t)$ is η -subgaussian. We let $\beta^* \in \mathcal{B}$ denote the parameters for the true decision objective. We evaluate parameter estimates $\beta \in \mathcal{B}$ using two metrics: (1) global prediction accuracy, and (2) the optimistic bias of the resulting policy π^β .

DEFINITION EC.1. Given $\beta \in \mathcal{B}$, the *excess mean-squared error (MSE)* is the generalization error relative to the true model (typically estimated on a held-out test set):

$$\mathcal{E}(\beta) = \mathbb{E}_{x,t} [(f^\beta(x, t) - f^{\beta^*}(x, t))^2].$$

DEFINITION EC.2. Given $\beta, \beta' \in \mathcal{B}$, the *optimism bias* of β when evaluated using β' is

$$\mathcal{R}(\beta; \beta') = \mathcal{P}(\beta; \beta') - \mathcal{P}(\beta; \beta^*) \quad \text{where} \quad \mathcal{P}(\beta; \beta') = \mathbb{E}_x [f^{\beta'}(x, \pi^\beta(x))].$$

To estimate β^* , we consider a historical dataset $W = \{(z_i, y_i)\}_{i=1}^n$, where $y_i = f^{\beta^*}(x_i, t_i) + \epsilon_i$ with i.i.d. noise $\epsilon_i \sim \mathcal{N}(0, \sigma^2)$; we let $\mathbb{P}_{x,t,\epsilon} = \mathbb{P}_{x,t} \times \mathbb{P}_\epsilon$ denote the distribution of a single historical example. We employ the classical ridge regression estimator, which serves as a proxy for any stable, regularized machine learning method:

$$\hat{\beta}^{(\lambda)} = \arg \min_{\beta \in \mathbb{R}^m} \left\{ \frac{1}{n} \|Z\beta - Y\|_2^2 + \lambda \|\beta\|_2^2 \right\},$$

where $\lambda \in \mathbb{R}_{>0}$ is the regularization hyperparameter, and

$$Z = \begin{bmatrix} - & z_1^\top & - \\ \vdots & \vdots & \vdots \\ - & z_n^\top & - \end{bmatrix}, \quad Y = \begin{bmatrix} y_1 \\ \vdots \\ y_n \end{bmatrix}.$$

Let $\Sigma = \mathbb{E}_{x,t}[zz^\top]$ be the true covariance matrix and $\hat{\Sigma} = n^{-1}Z^\top Z$ be the empirical covariance matrix, and let $\Sigma_\lambda = \Sigma + \lambda I$ and $\hat{\Sigma}_\lambda = \hat{\Sigma} + \lambda I$, where $I \in \mathbb{R}^{m \times m}$ is the identity matrix. Then,

$$\hat{\beta}^{(\lambda)} = (n\hat{\Sigma}_\lambda)^{-1}Z^\top Y = (I - \lambda\hat{\Sigma}_\lambda^{-1})\beta^* + \frac{1}{n}\hat{\Sigma}_\lambda^{-1}Z^\top E,$$

where $E = [\epsilon_1 \cdots \epsilon_n]^\top$ is the vector of noise terms. Given $\Delta \in \mathbb{R}_{>0}$, define the event

$$E_\Delta = \mathbb{1}(\|\hat{\Delta}\|_2 \leq \Delta) \quad \text{where} \quad \hat{\Delta} = \hat{\Sigma} - \Sigma.$$

In addition, given $\zeta \in \mathbb{R}_{>0}$ satisfying $\zeta \in \mathbb{R}_{>0}$, define the event

$$E'_\zeta = \mathbb{1}\left(\|\hat{\Sigma}_\lambda^{-1}Z^\top E\|_2^2 \leq \frac{5nm\sigma^2\zeta}{\lambda} \wedge \|\hat{\Sigma}_\lambda^{-1/2}Z^\top E\|_2^2 \leq 5nm\sigma^2\zeta\right).$$

We let $\mathbb{P}_E = (\mathbb{P}_\epsilon)^n$ denote the measure of E , and $\mathbb{P}_Z = (\phi_*\mathbb{P}_{x,t})^n$ of Z (where $f_*\beta$ denotes the pushforward measure).

EC.1.2. Results for the General Case

First, we prove that E_Δ holds with high probability as n grows large.

LEMMA EC.1. *We have*

$$\mathbb{P}_Z[E_\Delta] \geq 1 - 2m^2e^{-n\Delta^2/(2m^2\eta)^2}.$$

Proof. For every $i, j \in [m]$, by Hoeffding's inequality, $\mathbb{P}_Z[|\hat{\Delta}_{ij}| \geq \Delta/m] \leq 2e^{-n\Delta^2/(2m^2\eta^2)}$, so the claim follows by a union bound. Q.E.D.

Next, we prove a high-probability bound for E'_ζ .

LEMMA EC.2. *For $\zeta \geq 1$, we have*

$$\mathbb{P}_W[E'_\zeta] \geq 1 - e^{-\zeta}.$$

Proof. Let $Z/\sqrt{n} = U\hat{\Lambda}V^\top$ be the singular value decomposition of Z/\sqrt{n} , where

$$\hat{\Lambda} = \begin{bmatrix} \tilde{\Lambda} \\ 0_{(n-d) \times d} \end{bmatrix} \quad \text{where} \quad \tilde{\Lambda} = \begin{bmatrix} \lambda_1 & 0 & \cdots & 0 \\ 0 & \lambda_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \lambda_m \end{bmatrix}$$

for some $\lambda_1, \dots, \lambda_m \in \mathbb{R}_{>0}$, and $U \in \mathbb{R}^{n \times n}$ and $V \in \mathbb{R}^{m \times m}$ are orthogonal. Then, we have

$$\hat{\Sigma}_\lambda = \frac{1}{n}Z^\top Z + \lambda I = V\hat{\Lambda}^2V^\top + \lambda I = V(\tilde{\Lambda} + \lambda I)V^\top.$$

Thus,

$$\hat{\Sigma}_\lambda^{-1}Z^\top E = V(\tilde{\Lambda} + \lambda I)^{-1}\hat{\Lambda}UE.$$

Let $\tilde{E} = [\tilde{\epsilon}_1 \cdots \tilde{\epsilon}_m]$ with $\tilde{\epsilon}_i = (UE)_i$ for $i \in [m]$; note that since U is orthogonal, $UE \sim \mathcal{N}(0, \sigma^2 I_n)$ has the same distribution as E , so $\tilde{E} \sim \mathcal{N}(0, \sigma^2 I_m)$. Furthermore, we have

$$\hat{\Sigma}_\lambda^{-1} Z^\top E = \sqrt{n} V \tilde{\Lambda}' \tilde{E} \quad \text{where} \quad \tilde{\Lambda}' = \begin{bmatrix} \frac{\lambda_1}{\lambda_1^2 + \lambda} & 0 & \cdots & 0 \\ 0 & \frac{\lambda_2}{\lambda_2^2 + \lambda} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \frac{\lambda_m}{\lambda_m^2 + \lambda} \end{bmatrix}.$$

Now, note that

$$\|\hat{\Sigma}_\lambda^{-1} Z^\top E\|_2^2 = n \sum_{i=1}^m \frac{\lambda_i^2}{(\lambda_i^2 + \lambda)^2} \cdot \tilde{\epsilon}_i^2 = n \sum_{i=1}^m \left(\frac{\lambda_i^2 + \lambda}{(\lambda_i^2 + \lambda)^2} - \frac{\lambda}{(\lambda_i^2 + \lambda)^2} \right) \cdot \tilde{\epsilon}_i^2 \leq \frac{n}{\lambda} \sum_{i=1}^m \tilde{\epsilon}_i^2 = \frac{n\sigma^2}{\lambda} \sum_{i=1}^m \frac{\tilde{\epsilon}_i^2}{\sigma^2}.$$

Similarly, we have

$$\hat{\Sigma}_\lambda^{-1/2} Z^\top E = V(\tilde{\Lambda} + \lambda I)^{-1/2} \hat{\Lambda} U E = \sqrt{n} V \tilde{\Lambda}'' \tilde{E} \quad \text{where} \quad \tilde{\Lambda}'' = \begin{bmatrix} \frac{\lambda_1}{\sqrt{\lambda_1^2 + \lambda}} & 0 & \cdots & 0 \\ 0 & \frac{\lambda_2}{\sqrt{\lambda_2^2 + \lambda}} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \frac{\lambda_m}{\sqrt{\lambda_m^2 + \lambda}} \end{bmatrix}.$$

Now, note that

$$\|\hat{\Sigma}_\lambda^{-1/2} Z^\top E\|_2^2 = n \sum_{i=1}^m \frac{\lambda_i^2}{\lambda_i^2 + \lambda} \cdot \tilde{\epsilon}_i^2 \leq n \sum_{i=1}^m \tilde{\epsilon}_i^2 = n\sigma^2 \sum_{i=1}^m \frac{\tilde{\epsilon}_i^2}{\sigma^2}.$$

Finally, by the Laurent-Massart inequality, we have

$$\mathbb{P}_{\tilde{E}} \left[\sum_{i=1}^m \frac{\tilde{\epsilon}_i^2}{\sigma^2} \leq m + 2m\sqrt{\zeta} + 2\zeta \right] \geq 1 - e^{-\zeta},$$

so the claim follows. **Q.E.D.**

Next, we prove a standard formula for the excess MSE of any parameter vector β .

LEMMA EC.3. *We have $\mathcal{E}(\beta) = \|\Sigma^{1/2}(\beta - \beta^*)\|_2^2$.*

Proof. Note that

$$\begin{aligned} \mathcal{E}(\beta) &= \mathbb{E}_{x,t,\epsilon} [(\beta^\top z - y)^2] - \mathbb{E}[(\beta^{*\top} z - y)^2] \\ &= \mathbb{E}_{x,t,\epsilon} [(\beta^\top z - \beta^{*\top} z - \epsilon)^2] - \mathbb{E}_{x,t,\epsilon} [\epsilon^2] \\ &= \mathbb{E}_{x,t,\epsilon} [((\beta - \beta^*)^\top z)^2] \\ &= \mathbb{E}_{x,t,\epsilon} [\text{tr}((\beta - \beta^*)(\beta - \beta^*)^\top z z^\top)] \\ &= \text{tr}((\beta - \beta^*)(\beta - \beta^*)^\top \Sigma) \\ &= \|\Sigma^{1/2}(\beta - \beta^*)\|_2^2, \end{aligned}$$

as claimed. **Q.E.D.**

Our next result establishes a straightforward bound on the sensitivity of the estimated policy improvement $\mathcal{P}(\beta; \beta')$ to β' (note, however that \mathcal{P} can be arbitrarily sensitive to β).

LEMMA EC.4. *Given $\beta, \beta_1, \beta_2 \in \mathcal{B}$, if $\|\beta_1 - \beta_2\|_2 \leq \epsilon$, then $|\mathcal{P}(\beta; \beta_1) - \mathcal{P}(\beta; \beta_2)| \leq \epsilon \phi_{max}$.*

Proof. Note that

$$\begin{aligned} |\mathcal{P}(\beta; \beta_1) - \mathcal{P}(\beta; \beta_2)| &= |\mathbb{E}_x[(\beta_1 - \beta_2)^\top \phi(x, \pi^\beta(x))]| \\ &\leq \|\beta_1 - \beta_2\|_2 \cdot \mathbb{E}_x[\|\phi(x, \pi^\beta(x))\|_2] \\ &\leq \epsilon \phi_{max}, \end{aligned}$$

as claimed. Q.E.D.

Another simple result bounds the maximum eigenvalue of $\hat{\Sigma}_\lambda^{-1}$.

LEMMA EC.5. *We have $\|\hat{\Sigma}_\lambda^{-1}\|_2 \leq \lambda^{-1}$.*

Proof. Note that

$$\|\hat{\Sigma}_\lambda^{-1}\|_2 = \|(\hat{\Sigma} + \lambda I)^{-1}\|_2 = \frac{1}{\lambda_{\min}(\hat{\Sigma} + \lambda I)} = \frac{1}{\lambda + \lambda_{\min}(\hat{\Sigma})} \leq \frac{1}{\lambda},$$

where the last step follows since $\hat{\Sigma}$ is positive semi-definite. Q.E.D.

Next, we turn to analyzing the quantity $\hat{\Sigma}_\lambda^{-1}\Sigma$; this quantity is a key part of our analysis of the excess MSE of $\hat{\beta}^{(\lambda)}$.

LEMMA EC.6. *On event E_Δ with $\Delta \leq \lambda$, we have $\|\hat{\Sigma}_\lambda^{-1}\Sigma\|_2 \leq 1$.*

Proof. We have

$$\begin{aligned} \|\hat{\Sigma}_\lambda^{-1}\Sigma\|_2 &= \|(\Sigma + \hat{\Delta} + \lambda I)^{-1}\Sigma\|_2 \\ &= \|(I + \Sigma^{-1}(\hat{\Delta} + \lambda I))^{-1}\|_2 \\ &= \frac{1}{\lambda_{\min}(I + \Sigma^{-1}(\hat{\Delta} + \lambda I))} \\ &= \frac{1}{1 + \lambda_{\min}(\Sigma^{-1}(\hat{\Delta} + \lambda I))} \\ &\leq 1, \end{aligned}$$

where we have used the fact that $\Sigma^{-1}(\hat{\Delta} + \lambda I)$ is positive semi-definite on event E_Δ . Q.E.D.

Now, we state and prove our first main result, which bounds the accuracy of $\hat{\beta}^{(\lambda)}$ on events E_Δ and E'_ζ . Intuitively, the first term in the bound is the bias term (which becomes small as $\lambda \rightarrow 0$), and the second term is the variance term (which becomes small as $n \rightarrow \infty$). We generally think of λ as scaling as $1/\sqrt{n}$.

PROPOSITION EC.1 (Accuracy). *On event E_Δ with $\Delta \leq \lambda$ and E'_ζ with $\zeta \geq 1$, we have*

$$\mathcal{E}(\hat{\beta}^{(\lambda)}) \leq 2m\lambda \cdot \|\beta^*\|_2^2 + \frac{8m^2\sigma^2\zeta}{n}.$$

Proof. By Lemma EC.3, we have $\mathcal{E}(\hat{\beta}^{(\lambda)}) = \|\Sigma^{1/2}(\hat{\beta}^{(\lambda)} - \beta^*)\|_2^2$. Note that

$$\begin{aligned} \|\Sigma^{1/2}(\hat{\beta}^{(\lambda)} - \beta^*)\|_2 &= \frac{1}{n} \|\Sigma^{1/2} \hat{\Sigma}_\lambda^{-1} (-n\lambda\beta^* + Z^\top E)\|_2 \\ &\leq \lambda \|\Sigma^{1/2} \Sigma_\lambda^{-1}\|_2 \cdot \|\beta^*\|_2 + \frac{1}{n} \|\Sigma^{1/2} \hat{\Sigma}_\lambda^{-1} Z^\top E\|_2 \\ &\leq \lambda \|\beta^*\|_2 \cdot \sqrt{m \|\Sigma_\lambda^{-1} \Sigma\|_2 \cdot \|\Sigma_\lambda^{-1}\|_2} + \frac{1}{n} \|(\hat{\Sigma}_\lambda - \hat{\Delta} - \lambda I)^{1/2} \hat{\Sigma}_\lambda^{-1} Z^\top E\|_2 \\ &\leq \sqrt{m\lambda} \cdot \|\beta^*\|_2 + \frac{1}{n} \sqrt{\text{tr} \left(E^\top Z \hat{\Sigma}_\lambda^{-1} (\hat{\Sigma}_\lambda - \hat{\Delta} - \lambda I) \hat{\Sigma}_\lambda^{-1} Z^\top E \right)} \\ &\leq \sqrt{m\lambda} \cdot \|\beta^*\|_2 + \frac{1}{n} \sqrt{\|\hat{\Sigma}_\lambda^{-1/2} Z^\top E\|_2^2 + m \|\hat{\Delta} + \lambda I\|_2 \cdot \|\hat{\Sigma}_\lambda^{-1} Z^\top E\|_2^2} \\ &\leq \sqrt{m\lambda} \cdot \|\beta^*\|_2 + \sqrt{\frac{5m\sigma^2\zeta}{n} \left(1 + \frac{m\Delta + \lambda}{\lambda} \right)} \\ &\leq \sqrt{m\lambda} \cdot \|\beta^*\|_2 + 4m\sigma \sqrt{\frac{\zeta}{n}}, \end{aligned}$$

where the second-to-last step follows by Lemmas EC.5 & EC.6 and events E_Δ and E'_ζ . The claim follows from the inequality $(a+b)^2 \leq 2(a^2+b^2)$. **Q.E.D.**

Our second major result provides a stability guarantee for $\hat{\beta}^{(\lambda)}$; in particular, it says $\hat{\beta}^{(\lambda)}$ concentrates around a deterministic quantity $\bar{\beta}^{(\lambda)}$ on events E_Δ and E'_ζ (which, by Lemmas EC.1 & EC.2, holds with high probability over W). This result straightforwardly implies that all $\hat{\beta}^{(\lambda)}$ are pairwise close together. One thing to note in this result is that both Δ and λ scale as $1/\sqrt{n}$; thus, to obtain a bound that goes to zero, we need λ/Δ to be sufficiently large.

PROPOSITION EC.2 (Stability). *Letting $\bar{\beta}^{(\lambda)} = (I - \lambda\Sigma_\lambda^{-1})\beta^*$, on event E_Δ with $\Delta \leq \lambda/4$ and event E'_ζ with $\zeta \in \mathbb{R}_{>0}$, we have*

$$\|\hat{\beta}^{(\lambda)} - \bar{\beta}^{(\lambda)}\|_2^2 \leq \frac{4\Delta^2 \|\beta^*\|_2^2}{\lambda^2} + \frac{5m\sigma^2\zeta}{n\lambda}.$$

Proof. We prove a slightly more general result, where we let $\bar{\beta}^{(\lambda)} = (I - \lambda(\Sigma + \Gamma + \lambda I)^{-1})\beta^*$ for some positive semi-definite $\Gamma \in \mathbb{R}^{m \times m}$ satisfying $\gamma = \|\Gamma\|_2 \leq \lambda/4$. Then, letting $\hat{\Gamma} = \hat{\Delta} - \Gamma$ (so $\|\hat{\Gamma}\|_2 \leq \gamma + \Delta$), we show that

$$\|\hat{\beta}^{(\lambda)} - \bar{\beta}^{(\lambda)}\|_2^2 \leq \frac{4(\Delta + \gamma)^2 \|\beta^*\|_2^2}{\lambda^2} + \frac{5m\sigma^2\zeta}{n\lambda}.$$

The original result follows from taking $\Gamma = 0$, so $\bar{\beta}^{(\lambda)}$ is unchanged and $\gamma = 0$. Now, note that

$$\|\hat{\beta}^{(\lambda)} - \bar{\beta}^{(\lambda)}\|_2^2 \leq \lambda^2 \|\beta^*\|_2^2 \cdot \left\| (\Sigma + \hat{\Delta} + \lambda I)^{-1} - (\Sigma + \Gamma + \lambda I)^{-1} \right\|_2^2 + \frac{1}{n^2} \|\hat{\Sigma}_\lambda^{-1} Z^\top E\|_2^2.$$

For the first term, we have

$$\begin{aligned}
\|(\Sigma + \hat{\Delta} + \lambda I)^{-1} - (\Sigma + \Gamma + \lambda I)^{-1}\|_2 &= \left\| \left((I + (\Sigma + \Gamma + \lambda I)^{-1} \hat{\Gamma})^{-1} - I \right) (\Sigma + \Gamma + \lambda I)^{-1} \right\|_2 \\
&\leq \frac{1}{\lambda} \cdot \| (I + (\Sigma + \Gamma + \lambda I)^{-1} \hat{\Gamma})^{-1} - I \|_2 \\
&= \frac{1}{\lambda} \cdot \left\| \sum_{i=1}^{\infty} (-(\Sigma + \Gamma + \lambda I)^{-1} \hat{\Gamma})^i \right\|_2 \\
&\leq \frac{1}{\lambda} \sum_{i=1}^{\infty} \| (\Sigma + \Gamma + \lambda I)^{-1} \hat{\Gamma} \|_2^i \\
&\leq \frac{1}{\lambda} \sum_{i=1}^{\infty} \left(\frac{\Delta + \gamma}{\lambda} \right)^i \\
&= \frac{\Delta + \gamma}{\lambda^2 (1 - (\Delta + \gamma)/\lambda)} \\
&\leq \frac{2(\Delta + \gamma)}{\lambda^2},
\end{aligned}$$

where the Taylor expansion is justified since $\|(\Sigma + \Gamma + \lambda I)^{-1} \hat{\Gamma}\|_2 < 1$. The bound on the second term follows immediately from event E'_ζ . Q.E.D.

Our next result is a straightforward consequence of our previous result—it says that since $\hat{\beta}^{(\lambda)}$ concentrates to $\bar{\beta}^{(\lambda)}$, then the optimism bias when evaluating under $\hat{\beta}^{(\lambda)}$ correspondingly concentrates to the optimism bias under $\bar{\beta}^{(\lambda)}$.

PROPOSITION EC.3. *For any $\beta \in \mathbb{R}^m$, letting $\bar{\beta}^{(\lambda)} = (I - \lambda \Sigma_\lambda^{-1}) \beta^*$, on event E_Δ for some $\Delta \leq \lambda/4$ and on event E'_ζ for some $\zeta \in \mathbb{R}_{>0}$, we have*

$$|\mathcal{R}(\beta; \hat{\beta}^{(\lambda)}) - \mathcal{R}(\beta; \bar{\beta}^{(\lambda)})| = |\mathcal{P}(\beta; \hat{\beta}^{(\lambda)}) - \mathcal{P}(\beta; \bar{\beta}^{(\lambda)})| \leq \phi_{\max} \cdot \sqrt{\frac{4\Delta^2 \|\beta^*\|_2^2}{\lambda^2} + \frac{5m\sigma^2\zeta}{n\lambda}}.$$

Proof. We prove a slightly more general result, where we let $\bar{\beta}^{(\lambda)} = (I - \lambda(\Sigma + \Gamma + \lambda I)^{-1}) \beta^*$ for some positive semi-definite $\Gamma \in \mathbb{R}^{m \times m}$ satisfying $\gamma = \|\Gamma\|_2 \leq \lambda/4$. Then, letting $\hat{\Gamma} = \hat{\Delta} - \Gamma$ (so $\|\hat{\Gamma}\|_2 \leq \Delta + \gamma$), we show that

$$|\mathcal{R}(\beta; \hat{\beta}^{(\lambda)}) - \mathcal{R}(\beta; \bar{\beta}^{(\lambda)})| = |\mathcal{P}(\beta; \hat{\beta}^{(\lambda)}) - \mathcal{P}(\beta; \bar{\beta}^{(\lambda)})| \leq \phi_{\max} \cdot \sqrt{\frac{4(\Delta + \gamma)^2 \|\beta^*\|_2^2}{\lambda^2} + \frac{5m\sigma^2\zeta}{n\lambda}}.$$

The original result follows from taking $\Gamma = 0$, so $\bar{\beta}^{(\lambda)}$ is unchanged and $\gamma = \Delta$. Now, note that

$$|\mathcal{R}(\beta; \hat{\beta}^{(\lambda)}) - \mathcal{R}(\beta; \bar{\beta}^{(\lambda)})| = |\mathcal{P}(\beta; \hat{\beta}^{(\lambda)}) - \mathcal{P}(\beta; \bar{\beta}^{(\lambda)})| \leq \|\hat{\beta}^{(\lambda)} - \bar{\beta}^{(\lambda)}\|_2 \cdot \phi_{\max},$$

where the inequality follows by Lemma EC.4. Then, the result follows from Proposition EC.2.

EC.1.3. Proof of Proposition 3

Let $\Delta = \sqrt{18 \log(8/\delta)/n}$ (so $\lambda = \alpha\Delta$) and $\zeta = \log(8/\delta)$. Note that $\|\beta^*\|_2 = b$, and

$$\phi_{\max} = \mathbb{E}_x \left[\max_{t \in \mathcal{T}} \sqrt{x^2 + t^2 + (xt)^2} \right] = \mathbb{E}_x [\sqrt{2x^2 + 1}] \leq \sqrt{2} \cdot \mathbb{E}_x [|x|] + 1 = \frac{2}{\sqrt{\pi}} + 1 \leq 3.$$

Next, by Lemma EC.1, we have $\mathbb{P}_Z[E_\Delta] \geq 1 - \delta/8$, and by Lemma EC.2, we have $\mathbb{P}_W[E'_\zeta] \geq 1 - \delta/8$. By a union bound, these events hold for all $\beta \in \{\hat{\beta}^{(\lambda)}, \hat{\beta}^{(\lambda)'}, \hat{\beta}^{(\lambda)''}, \hat{\beta}^{(\lambda)'''}\}$. Now, we prove the three results on the events E_Δ and E'_ζ for all of these β .

Accuracy. By Proposition EC.1, on event E_Δ and E'_ζ , we have

$$\mathcal{E}(\hat{\beta}^{(\lambda)}) \leq 2m\lambda \cdot \|\beta^*\|_2^2 + \frac{8m^2\sigma^2\zeta}{n} = 6b^2\alpha \sqrt{\frac{18 \log(8/\delta)}{n}} + \frac{72 \log(8/\delta)}{n}.$$

Stability. Let

$$\bar{\Sigma} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 1 \\ 0 & 1 & 1 \end{bmatrix} \quad \text{and} \quad \Gamma = \begin{bmatrix} 1-p & 0 & 0 \\ 0 & 0 & 1-p \\ 0 & 1-p & 1-p \end{bmatrix},$$

so $\bar{\Sigma} = \Sigma + \Gamma$. Further define $\bar{\Sigma}_\lambda = \bar{\Sigma} + \lambda I$. Also, note that $\gamma = \|\Gamma\|_2 = 2(1-p) = \Delta$. Furthermore, let $\bar{\beta}^{(\lambda)} = (I - \lambda \bar{\Sigma}_\lambda^{-1})\beta^*$. Then, on events E_Δ and E'_ζ , by Proposition EC.2, we have

$$\|\hat{\beta}^{(\lambda)} - \bar{\beta}^{(\lambda)}\|_2 \leq \sqrt{\frac{4(\Delta + \gamma)^2 \|\beta^*\|_2^2}{\lambda^2} + \frac{5m\sigma^2\zeta}{n\lambda}} \leq \sqrt{\frac{4b^2(\Delta + \gamma)^2}{\lambda^2} + \frac{15 \log(8/\delta)}{n\lambda}} \leq \frac{4b}{\alpha} + 2\sqrt{\frac{\log(8/\delta)}{\alpha\sqrt{n}}},$$

so the bound on $\|\hat{\beta}^{(\lambda)} - \hat{\beta}^{(\lambda)'}\|_2$ follows by the triangle inequality. Next, by Proposition EC.3, we have

$$|\mathcal{P}(\hat{\beta}^{(\lambda)}, \hat{\beta}^{(\lambda)'}) - \mathcal{P}(\hat{\beta}^{(\lambda)}; \bar{\beta}^{(\lambda)})| \leq \phi_{\max} \|\hat{\beta}^{(\lambda)} - \bar{\beta}^{(\lambda)}\|_2 \leq \frac{12b}{\alpha} + 6\sqrt{\frac{\log(8/\delta)}{\alpha\sqrt{n}}}.$$

Now, we bound $|\mathcal{P}(\hat{\beta}^{(\lambda)}; \bar{\beta}^{(\lambda)}) - \mathcal{P}(\bar{\beta}^{(\lambda)}; \bar{\beta}^{(\lambda)})|$. To this end, let $\hat{\xi} = \hat{\beta}^{(\lambda)} - \bar{\beta}^{(\lambda)}$, and note that

$$\begin{aligned} \pi^{\bar{\beta}^{(\lambda)}}(x) &= \mathbb{1}(x \geq 0) \\ \pi^{\hat{\beta}^{(\lambda)}}(x) &= \mathbb{1}\left(x \geq \frac{-\hat{\xi}_1}{((b/2) + \hat{\xi}_3)}\right). \end{aligned}$$

On events E_Δ and E'_ζ , we have $\|\hat{\xi}\|_2 \leq \xi_{\max}$, where

$$\xi_{\max} = \frac{4b}{\alpha} + 2\sqrt{\frac{\log(8/\delta)}{\alpha\sqrt{n}}}.$$

By our assumptions on n and α , we have $\xi_{\max} \leq b/4$. Thus, defining $I = [-4\xi_{\max}/b, 4\xi_{\max}/b]$, then

$\pi^{\hat{\beta}^{(\lambda)}}(x) = \pi^{\bar{\beta}^{(\lambda)}}(x)$ for $x \notin I$. As a consequence, we have

$$|\mathcal{P}(\hat{\beta}^{(\lambda)}; \bar{\beta}^{(\lambda)}) - \mathcal{P}(\bar{\beta}^{(\lambda)}; \bar{\beta}^{(\lambda)})| = \left| \mathbb{E}_x \left[\frac{x(\pi^{\hat{\beta}^{(\lambda)}}(x) - \pi^{\bar{\beta}^{(\lambda)}}(x))}{2} \right] \right| \leq \frac{2\xi_{\max}}{b} \cdot \mathbb{P}_x[x \in I] \leq \frac{2\xi_{\max}}{b} \cdot \frac{|I|}{2\sqrt{2\pi}} = \frac{8\xi_{\max}^2}{b^2\sqrt{2\pi}}.$$

By the inequality $(a + b)^2 \leq 2a^2 + 2b^2$, we have

$$\frac{8\xi_{\max}^2}{b^2\sqrt{2\pi}} \leq \frac{128}{\alpha^2} + \frac{32\log(8/\delta)}{\alpha b^2\sqrt{n}}.$$

Finally, by repeated application of the triangle inequality, we have

$$|\mathcal{P}(\hat{\beta}^{(\lambda)}; \hat{\beta}^{(\lambda)'}) - \mathcal{P}(\hat{\beta}^{(\lambda)''}; \hat{\beta}^{(\lambda)'''})| \leq \frac{24b}{\alpha} + 12\sqrt{\frac{\log(8/\delta)}{\alpha\sqrt{n}}} + \frac{128}{\alpha^2} + \frac{32\log(8/\delta)}{\alpha b^2\sqrt{n}}.$$

Winner's curse. Note that

$$\bar{\Sigma}_\lambda^{-1} = \begin{bmatrix} \frac{1}{1+\lambda} & 0 & 0 \\ 0 & \frac{1+\lambda}{\lambda(2+\lambda)} & -\frac{1}{\lambda(2+\lambda)} \\ 0 & -\frac{1}{\lambda(2+\lambda)} & \frac{1+\lambda}{\lambda(2+\lambda)} \end{bmatrix},$$

so

$$\bar{\beta}^{(\lambda)} = (I - \lambda\bar{\Sigma}_\lambda^{-1})\beta^* = \begin{bmatrix} 1 - \frac{\lambda}{1+\lambda} & 0 & 0 \\ 0 & 1 - \frac{1+\lambda}{2+\lambda} & \frac{1}{2+\lambda} \\ 0 & \frac{1}{2+\lambda} & 1 - \frac{1+\lambda}{2+\lambda} \end{bmatrix} \begin{bmatrix} 0 \\ b \\ 0 \end{bmatrix} = \begin{bmatrix} 0 \\ \frac{b}{2+\lambda} \\ 0 \end{bmatrix} = \begin{bmatrix} 0 \\ \frac{b}{2+(1/n)} \\ \frac{b}{2+(1/n)} \end{bmatrix}.$$

Finally, we have $\pi^{\bar{\beta}^{(\lambda)}}(x) = \mathbb{1}(x \geq 0)$, so

$$\begin{aligned} \mathcal{R}(\bar{\beta}^{(\lambda)}; \bar{\beta}^{(\lambda)}) &= \mathbb{E}_x \left[\frac{bx}{2+(1/n)} + \frac{x\mathbb{1}(bx \geq 0)}{2+(1/n)} \right] - \mathbb{E} \left[\frac{bx}{2+(1/n)} + \frac{bt}{2+(1/n)} \right] \\ &= \frac{b}{2+(1/n)} \cdot \mathbb{E}[x\mathbb{1}(x \geq 0)] \\ &= \frac{b}{(2+(1/n))\sqrt{2\pi}} \\ &\geq \frac{b}{2\sqrt{2\pi}}. \end{aligned}$$

The claim follows by the stability of $\mathcal{P}(\hat{\beta}^{(\lambda)}; \hat{\beta}^{(\lambda)'})$ established above together with the fact that by our choice of β^* , for all $\beta \in \mathbb{R}^m$, we have $\mathcal{P}(\beta; \beta^*) = 0$, so $\mathcal{R}(\beta; \beta') = \mathcal{P}(\beta; \beta')$. Q.E.D.

EC.2. Details of Refugee Matching Case Study

This section provides the implementation details for our refugee matching simulation environment. Wherever possible, we calibrated parameters, sample sizes, and marginal distributions using those reported for the US site in the Supplemental Material of Bansak et al. (2018).

EC.2.1. Additional Details on Bansak et al. (2018) and Ahani et al. (2021)

Recent work proposes to use historical data to design algorithmic assignment rules that improve short-run employment relative to existing practice.

Using U.S. and Swiss data, Bansak et al. (2018) train machine learning models to predict employment for each refugee–location pair using features such as country of origin, language, age, gender, education, and placement restrictions. In their main specification, they fit separate prediction models by location and select gradient-boosted trees based on out-of-sample classification accuracy and calibration. They then:

1. estimate predicted employment probabilities for each refugee at each feasible location;
2. aggregate individual predictions to a family-level score (for example, the predicted probability that at least one family member is employed at a location); and
3. solve a linear assignment problem to maximize the average family-level score, subject to capacity and practical constraints.

To evaluate the algorithm, they run backtests: train on earlier arrivals, treat later arrivals with no placement restrictions as a test set, and compare predicted employment under the algorithmic assignment to the observed employment rate under the historical assignment. Reported gains are large, between 40–75% improvements in predicted employment relative to existing procedures.

Similarly, Ahani et al. (2021) develop Annie MOORE, an integrated machine learning and integer-optimization tool for a U.S. resettlement agency. They focus on “free cases”—families without U.S. ties who can be assigned flexibly across affiliates. Using administrative data on demographics, household structure, origin, language, health indicators, and macroeconomic conditions, they train predictive models of employment for refugee-affiliate pairs.

They compare pooled logistic regression, affiliate-specific logistic regression, LASSO-logit with hand-specified interactions, and gradient-boosted trees. In their test set, LASSO and gradient-boosted trees achieve substantially lower misclassification error and higher AUC than simpler models; LASSO is chosen as the primary model because it combines good discrimination with reasonable calibration.

Given these predictions, they:

1. treat the predicted employment probabilities as quality scores for each refugee–affiliate pair;
2. solve a mixed-integer program that allocates refugees to affiliates, maximizing total predicted employment subject to capacity and operational constraints; and
3. evaluate gains by comparing the sum of predicted employment under the optimized allocation to the observed employment under historical assignment, using the same LASSO model to impute counterfactual outcomes.

These backtests suggest improvements in predicted employment on the order of 20–40%. In an e-companion, they also conduct a bootstrap analysis: they draw many bootstrap samples of the training data, re-estimate the LASSO model on each, and evaluate the *same* optimized allocation under each bootstrapped model. The resulting distribution of estimated gains is tight and positive, which they interpret as evidence of robustness.

EC.2.2. Covariate Marginals and Location Assignment Probabilities

Covariate	Probability
Age \sim Unif(18, 30)	0.44
Age \sim Unif(30, 40)	0.28
Age \sim Unif(40, 50)	0.16
Age \sim Unif(50, 60)	0.12
Gender = Male	0.53
Education = None	0.18
Education = Little	0.39
Education = Secondary	0.21
Education = Advanced	0.10
Education = University	0.12
Speaks English = True	0.43
Case Restriction = Free	0.28
Origin = Burma	0.23
Origin = Iraq	0.2
Origin = Bhutan	0.13
Origin = Somalia	0.11
Origin = Afghanistan	0.07
Origin = Dem. Rep. of Congo	26/900
Origin = Iran	26/900
Origin = Eritrea	26/900
Origin = Ukraine	26/900
Origin = Syria	26/900
Origin = Sudan	26/900
Origin = Ethiopia	26/900
Origin = Moldova	26/900
Origin = Other	26/900
Arrival Year \sim Unif{2011, ..., 2015}	20/23
Arrival Year = 2016	3/23
Arrival Month \sim Unif{Jan, ..., Sep}	6/69
Arrival Month \sim Unif{Oct, Nov, Dec}	5/69

Table EC.1 Probability of covariates being assigned (independently) to each refugee, following the marginals reported in the Supplemental Material of Bansak et al. (2018). For binary covariates only one option is reported.

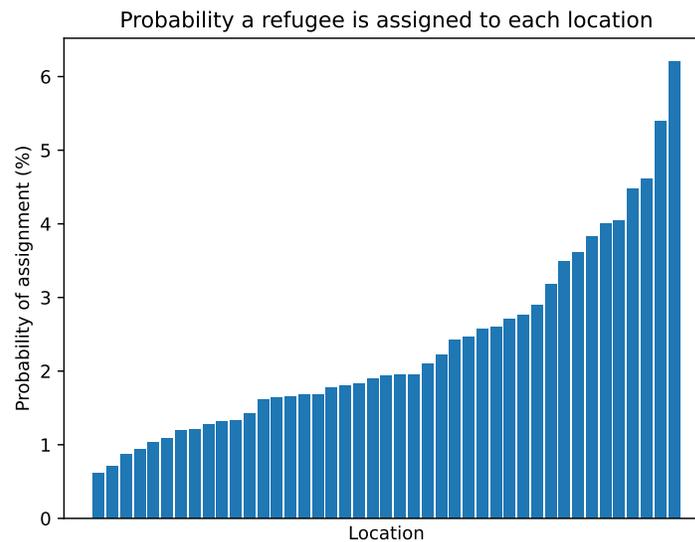


Figure EC.1 Probability of each location being assigned to each refugee sample, sorted in ascending order. This distribution follows the empirical distribution reported in Bansak et al. (2018).

EC.2.3. Employment Rates by Location

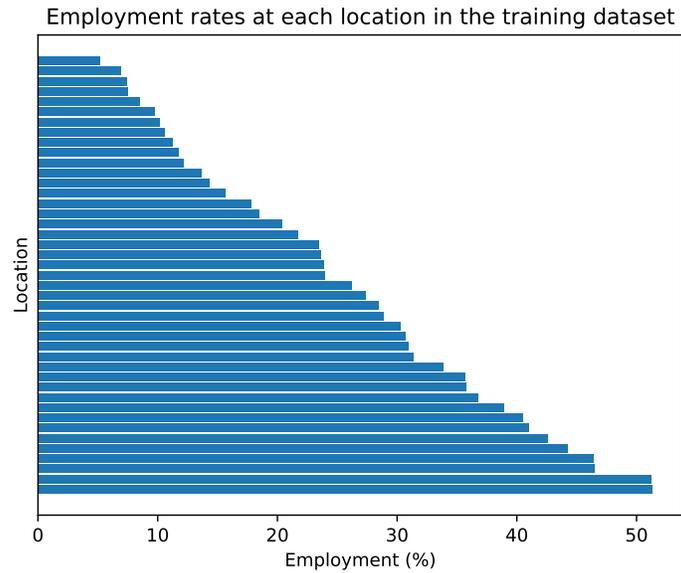


Figure EC.2 The distribution of employment rates across location have been tuned to resemble those in Bansak et al. (2018) to maximize similarities between the simulation environment and real world data.

EC.2.4. Prediction Model Calibration and ROC Curves

The calibration and ROC curves reflect that each prediction model is well-calibrated and makes fairly accurate predictions.

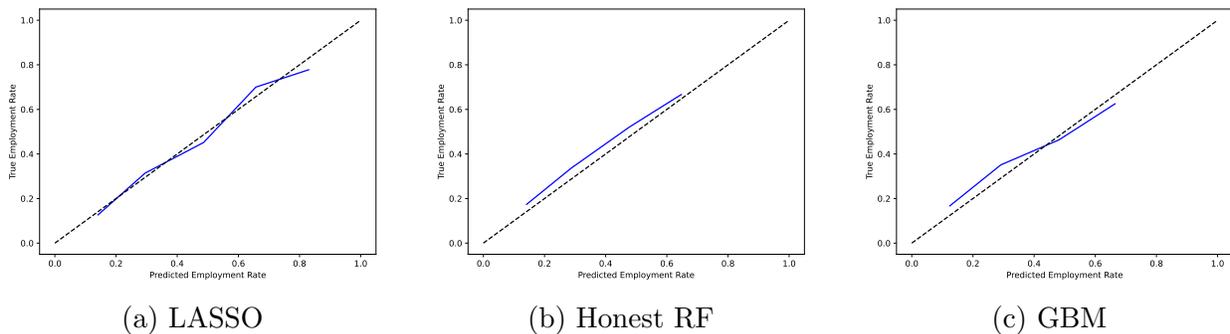
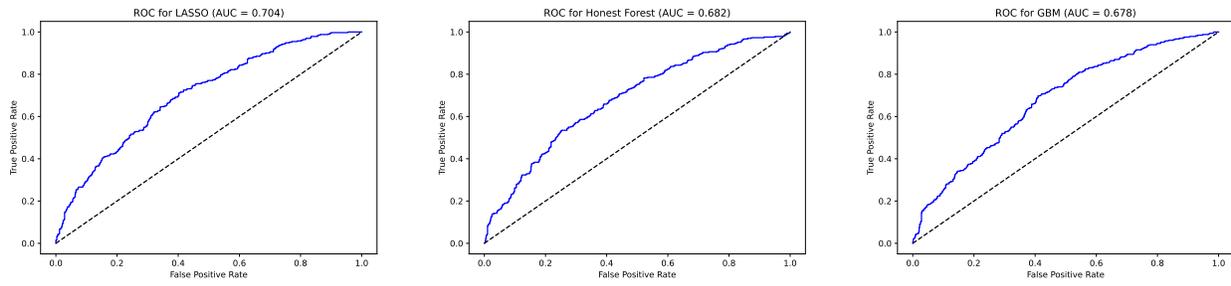


Figure EC.3 Calibration curves for each prediction model.



(a) LASSO (AUC = 0.704)

(b) Honest RF (AUC = 0.682)

(c) GBM (AUC = 0.678)

Figure EC.4 ROC curves detailing the performance of each prediction model on the test dataset.