

Appendix

A Organization of the Appendix

The Appendix is organized as follows.

- Section [A](#) provides the organization of the appendix.
- Section [B](#) provides the proof of Theorem [1](#).
- Section [C](#) includes notations and proofs for the discussion in section [3.3](#).
- Section [D](#) includes the proofs and derivations for section [4](#).
- Section [E](#) presents additional related works.
- Section [F](#) shows additional experimental details and results, including basic information on each dataset and the computing infrastructure.

B Proof of Theorem [1](#)

In this section, we provide the proof of Theorem [1](#). To simplify our discussion, we focus on the unconstrained best response, i.e. the case in which $F = A$. The proofs for the other two types of best response ($F = M$, $F = I$) follow the same arguments except that the inverse of $(S^{-1})_I$ does not equal to S , but equals to $((S^{-1})_I)^{-1}$.

We first prove two lemmas that allow us to reformulate the best response as an optimization problem. The first states that the decision subject’s goal is to maximize their utility, but they are unwilling to pay a cost greater than 2:

Lemma 1 (Decision Subject’s Best-Response Function). *Given a classifier $h : \mathcal{X} \rightarrow \{-1, +1\}$, a cost function $c : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$, and a set of realizable feature vectors $\mathcal{X}^\dagger \subseteq \mathcal{X}$, the best response of a decision subject with features $x \in \mathcal{X}^\dagger$ is the solution to the following optimization program:*

$$\max_{x' \in \mathcal{X}^\dagger} U(x, x') \quad \text{s.t.} \quad c(x, x') \leq 2$$

Proof. Since the classifier in our game outputs a binary decision (-1 or $+1$), decision subjects only have an incentive to change their features from x to x' when $c(x, x') \leq 2$. To see this, notice that a decision subject originally classified as -1 receives a default utility of $U(x, x) = f(x) - 0 = -1$ by presenting her original features x . Since costs are always non-negative, she can only hope to increase her utility by flipping the classifier’s decision. If she changes her features to some x' such that $f(x') = +1$, then the new utility will be given by

$$U(x, x') = f(x') - c(x, x') = 1 - c(x, x')$$

Hence the decision subject will only change her features if $1 - c(x, x') \geq f(x) = -1$, or $c(x, x') \leq 2$. \square

The next lemma turns the above maximization program into a minimization program, in which the decision subject seeks the minimum-cost change in x that crosses the decision boundary. If the cost exceeds 2, which is the maximum possible gain from adaptation, they would rather not modify any features.

Lemma 2. *Let x^* be an optimal solution to the following optimization problem:*

$$\begin{aligned} x^* &= \arg \min_{x' \in \mathcal{X}_A^*(x)} c(x, x') \\ \text{s.t.} \quad & \text{sign}(w^\top x') = 1 \end{aligned}$$

If no solution is returned, we say an x^ such that $c(x, x^*) = \infty$ is returned. Define $\Delta(x)$ as follows:*

$$\Delta(x) = \begin{cases} x^*, & \text{if } c(x, x^*) \leq 2 \\ x, & \text{otherwise} \end{cases}$$

Then $\Delta(x)$ is an optimal solution to the optimization problem in Lemma [1](#).

Proof. Recall that the utility function of the decision subject is $U(x, x') = f(x') - c(x, x')$, and that, by Lemma [1](#) they will only modify their features if the utility increases, i.e. if they achieve $f(x') = +1$ and while incurring cost $c(x, x') \leq 2$.

Consider two cases for $x' \neq x$:

1. When $c(x, x') > 2$, there are no feasible points for the optimization problem of Lemma [1](#).
2. When $c(x, x') \leq 2$, we only need to consider those feature vectors x' that satisfy $f(x') = 1$, because if $f(x') = -1$, the decision subject with features x would prefer not to change anything. Since maximizing $U(x, x') = f(x') - c(x, x')$ is equivalent to minimizing $c(x, x')$ if $f(x') = 1$, we conclude that when $c(x, x') \leq 2$, the optimum of the program of Lemma [1](#) is the same as the optimum of the program in Lemma [2](#).

□

Lemma [2](#) enables us to re-formulate the objective function as follows. Recall that $c(x, x') = \sqrt{(x_A - x_A')^\top S^{-1} (x_A - x_A')}$ where S^{-1} is symmetric positive definite. Thus S^{-1} has the following diagonalized form, in which Q is an orthogonal matrix and Λ^{-1} is a diagonal matrix:

$$S^{-1} = Q^\top \Lambda^{-1} Q = (\Lambda^{-\frac{1}{2}} Q)^\top (\Lambda^{-\frac{1}{2}} Q)$$

With this, we can re-write the cost function as

$$\begin{aligned} c(x, x') &= \sqrt{(x_A - x_A')^\top S^{-1} (x_A - x_A')} \\ &= \sqrt{(x_A - x_A')^\top (\Lambda^{-\frac{1}{2}} Q)^\top (\Lambda^{-\frac{1}{2}} Q) (x_A - x_A')} \\ &= \sqrt{(\Lambda^{-\frac{1}{2}} Q (x_A - x_A'))^\top (\Lambda^{-\frac{1}{2}} Q (x_A - x_A'))} \\ &= \|\Lambda^{-\frac{1}{2}} Q (x_A - x_A')\|_2 \end{aligned}$$

Meanwhile, the constraint in Lemma [2](#) can be written

$$\begin{aligned} \text{sign}(w \cdot x') &= \text{sign}(w_A \cdot x_A' + w_{\text{IM}} \cdot x_{\text{IM}}) \\ &= \text{sign}(w_A \cdot x_A' - (-w_{\text{IM}} \cdot x_{\text{IM}})) = 1 \end{aligned}$$

Hence the optimization problem can be reformulated as

$$\min_{x_A' \in \mathcal{X}_A^*} \|(\Lambda^{-\frac{1}{2}} Q (x_A - x_A'))\|_2 \tag{7}$$

$$\text{s.t. } \text{sign}(w_A \cdot x_A' - (-w_{\text{IM}} \cdot x_{\text{IM}})) = 1 \tag{8}$$

The above optimization problem can be further simplified by getting rid of the $\text{sign}(\cdot)$:

Lemma 3. *If x_A^\mp is an optimal solution to Eq. [\(7\)](#) under constraint Eq. [\(8\)](#), then it must satisfy $w_A \cdot x_A^\mp - (-w_{\text{IM}} \cdot x_{\text{IM}}) = 0$.*

Proof. We prove by contradiction. Let x_A^\mp is an optimal solution to Eq. [\(7\)](#) and suppose towards contradiction that $w_A \cdot x_A^\mp > -w_{\text{IM}} \cdot x_{\text{IM}}$. Since the original feature vector x was classified as -1 , we have

$$w_A \cdot x_A^\mp > -w_{\text{IM}} \cdot x_{\text{IM}}, \quad w_A \cdot x_A < -w_{\text{IM}} \cdot x_{\text{IM}}$$

By the continuity properties of linear vector space, there exists $\mu \in (0, 1)$ such that:

$$w_A \left(\mu \cdot x_A^\mp + (1 - \mu)x_A \right) = -w_{IM} \cdot x_{IM}$$

Let $x_A'' = \mu \cdot x_A^\mp + (1 - \mu)x_A$. Then $\text{sign}(w_A x_A'' - (-w_{IM} \cdot x_{IM})) = 1$, i.e., x_A'' also satisfies the constraint. Since x_A^\mp is an optimum of Eq. (7), we have

$$\|\Sigma^{-\frac{1}{2}}Q(x_A^\mp - x_A)\| \leq \|\Sigma^{-\frac{1}{2}}Q(x_A'' - x_A)\|$$

However, we also have:

$$\begin{aligned} \|\Sigma^{-\frac{1}{2}}Q(x_A'' - x_A)\| &= \|\Sigma^{-\frac{1}{2}}Q(\mu \cdot x_A^\mp + (1 - \mu)x_A - x_A)\| \\ &= \|\Sigma^{-\frac{1}{2}}Q(\mu \cdot (x_A^\mp - x_A))\| \\ &= \mu \|\Sigma^{-\frac{1}{2}}Q(x_A^\mp - x_A)\| \\ &< \|\Sigma^{-\frac{1}{2}}Q(x_A^\mp - x_A)\| \end{aligned}$$

contradicting our assumption that x_A^\mp is optimal. Therefore x_A^\mp must satisfy $w_A x_A^\mp = -w_{IM} \cdot x_{IM}$. \square

As a result of Lemma 3, we can replace the constraint in Eq. (7) with its corresponding equality constraint without changing the optimal solution.² The decision subject's best-response program from Lemma 1 is therefore equivalent to

$$\min_{x_A' \in \mathcal{X}_A^*} \|(\Lambda^{-\frac{1}{2}}Q(x_A - x_A'))\|_2 \quad (9)$$

$$\text{s.t. } w_A \cdot x_A' - (-w_{IM} \cdot x_{IM}) = 0 \quad (10)$$

The following lemma gives us a closed-form solution for the above optimization problem:

Lemma 4. *The optimal solution to the optimization problem defined in Eq. (9) and Eq. (10)*

has the following closed form:

$$x_A^\mp = x_A - \frac{w^\top x}{w_A^\top S w_A} S w_A.$$

Proof. Notice that the above program has the form

$$\begin{aligned} \min_{x_A' \in x_A^*} \|A x_A' - b\|_2 \\ \text{s.t. } C x_A' = d \end{aligned}$$

where $A = \Lambda^{-\frac{1}{2}}Q$, $b = \Lambda^{-\frac{1}{2}}Q x_A$, $C = w_A^\top$, and $d = -w_{IM}^\top x_{IM}$. Note the following useful equalities:

$$\begin{aligned} A^\top A &= (\Lambda^{-\frac{1}{2}}Q)^\top \Lambda^{-\frac{1}{2}}Q = S^{-1} \\ (A^\top A)^{-1} &= S \\ A^\top b &= (\Lambda^{-\frac{1}{2}}Q)^\top \Lambda^{-\frac{1}{2}}Q x_A = S^{-1} x_A \end{aligned}$$

²A similar argument was made by Haghtalab et al. (2020) but here we provide a proof for a more general case, where the objective function is to minimize a weighted norm instead of simply $\|x_A - x_A'\|_2$.

The above is a norm minimization problem with equality constraints, whose optimum x_A^\mp has the following closed form [Boyd & Vandenberghe \(2004\)](#):

$$\begin{aligned} x_A^\mp &= (A^\top A)^{-1} \left(A^\top b - C^\top (C(A^\top A)^{-1} C^\top)^{-1} (C(A^\top A)^{-1} A^\top b - d) \right) \\ &= S \left(S^{-1} x_A - w_A (w_A^\top S w_A)^{-1} (w_A^\top S (S^{-1} x_A) - (-w_{\text{IM}}^\top x_{\text{IM}})) \right) \\ &= x_A - S \left(w_A (w_A^\top S w_A)^{-1} (w_A^\top x_A + w_{\text{IM}}^\top x_{\text{IM}}) \right) \\ &= x_A - \frac{w^\top x}{w_A^\top S w_A} S w_A \end{aligned}$$

□

We can now compute the cost incurred by an individual with features x who plays their best response x^\mp :

$$\begin{aligned} c(x, x^\mp) &= \sqrt{(x_A - x_A^\mp)^\top S^{-1} (x_A - x_A^\mp)} \\ &= \sqrt{\left(\frac{w^\top x}{w_A^\top S w_A} S w_A \right)^\top S^{-1} \left(\frac{w^\top x}{w_A^\top S w_A} S w_A \right)} \\ &= \frac{|w^\top x|}{\sqrt{w_A^\top S w_A}} \end{aligned}$$

Hence an decision subject who was classified as -1 with feature vector x has the unconstrained best response

$$\Delta(x) = \begin{cases} x, & \text{if } \frac{|w^\top x|}{\sqrt{w_A^\top S w_A}} \geq 2 \\ \left[x_A - \frac{w^\top x}{w_A^\top S w_A} S w_A \mid x_{\text{IM}} \right], & \text{otherwise} \end{cases}$$

which completes the proof of [Theorem 1](#)

C Proofs of Propositions in [Section 3.3](#)

Notation. We make use of the following additional notation:

- $v^{(i)}$ denotes the i -th element of a vector v
- For any $F \in \{\text{A}, \text{I}, \text{M}\}$, $\Delta^F \in \mathbb{R}^{d^F}$ denotes the vector containing only features of type F within the best response $\Delta(x)$.
- $\mathbf{0}$ denotes the vector whose elements are all 0
- $A \succ B$ indicates that matrix $A - B$ is positive definite
- e_i denotes the vector containing 1 in its i -th component and 0 elsewhere

C.1 Proof of [Proposition 1](#)

Proof. Let $w_M^{(m)} \neq 0$, and consider an decision subject with original features x who was classified as -1 . By [Theorem 1](#) the actionable sub-vector of x 's unconstrained best response is

$$\Delta^{\text{A}}(x) = \frac{w^\top x}{w_A^\top S w_A} S \cdot w_A = \frac{w^\top x}{w_A^\top S w_A} \begin{bmatrix} S_{\text{I}} & 0 \\ 0 & S_{\text{M}} \end{bmatrix} \begin{bmatrix} w_{\text{I}} \\ w_{\text{M}} \end{bmatrix} = \frac{w^\top x}{w_A^\top S w_A} \begin{bmatrix} S_{\text{I}} \cdot w_{\text{I}} \\ S_{\text{M}} \cdot w_{\text{M}} \end{bmatrix}$$

And in particular,

$$\Delta^M(x) = \frac{w^\top x}{w_A^\top S w_A} S_M \cdot w_M$$

Since x was initially classified as -1 , we have $w^\top x < 0$, which means $\frac{w^\top x}{w_A^\top S w_A} \neq 0$. For convenience, let $c = \frac{w^\top x}{w_A^\top S w_A}$. We have

$$\Delta^M(x) - x_M = c S_M w_M - x_M = S_M (c w_M - S_M^{-1} x_M)$$

Now examine the following:

$$\begin{aligned} (c w_M - S_M^{-1} x_M)^{(m)} &= c w_M^{(m)} - (S_M^{-1} x_M)^{(m)} \\ &= c w_M^{(m)} - \sum_{i=1}^{d_M} (S_M^{-1})^{(im)} x_M^{(m)} \end{aligned}$$

Recall that $c w_M^{(m)} \neq 0$. Hence if $\sum_{i=1}^{d_M} (S_M^{-1})^{(im)} = 0$, or if

$$x_M^{(m)} \neq \frac{c w_M^{(m)}}{\sum_{i=1}^{d_M} (S_M^{-1})^{(im)}},$$

then $(c w_M - S_M^{-1} x_M)^{(m)} \neq 0$, and therefore $c w_M - S_M^{-1} x_M \neq \mathbf{0}$. Since S_M is positive definite, it has full rank, which implies

$$\Delta^M(x) - x_M = S_M (c w_M - S_M^{-1} x_M) \neq \mathbf{0}$$

as required. With this, we have shown that when there exists a manipulated feature $x^{(m)}$ whose corresponding coefficient $w_A^{(m)} \neq 0$, the classifier is vulnerable to changes in the manipulated features by the vast majority of decision subjects. \square

C.2 Proof of Proposition 2

Proof. Consider a decision subject with features x such that $h(x) = -1$. Suppose x can flip this classification result by performing the improving best response $\Delta_1(x)$, which implies that the cost of that action is no greater than 2 for this decision subject. We therefore have:

$$2 \geq c(x, \Delta_1(x)) = \frac{|w^\top x|}{\sqrt{w_1^\top S_1 w_1}} > \frac{|w^\top x|}{\sqrt{w_1^\top S_1 w_1 + w_M^\top S_M w_M}} = \frac{|w^\top x|}{\sqrt{w_A^\top S w_A}} = c(x, \Delta(x))$$

where the strict inequality is due to the fact that $S_M \succ 0$ and $w_M \neq \mathbf{0}$. As we have shown that $c(x, \Delta(x)) < 2$, we conclude whenever an decision subject can successfully flip her decision by the improving best response, she can also achieve it by performing the unconstrained best response.

On the other hand, consider the case when the unconstrained best response of a decision subject with features x^* has cost exactly 2:

$$2 = c(x^*, \Delta(x^*)) = \frac{|w^\top x^*|}{\sqrt{w_A^\top S w_A}} = \frac{|w^\top x^*|}{\sqrt{w_1^\top S_1 w_1 + w_M^\top S_M w_M}} < \frac{|w^\top x^*|}{\sqrt{w_1^\top S_1 w_1}} = c(x^*, \Delta_1(x^*))$$

where the strict inequality is due to the fact that $S_M \succ 0$ and $w_M \neq \mathbf{0}$. As we have shown that $c(x^*, \Delta_1(x^*)) > 2$, we conclude that while the unconstrained best response is viable for this decision subject, the improving best response is not. \square

C.3 Proof of Proposition 4

Proof. Let the cost covariance matrices for groups Φ and Ψ be

$$S_{\Psi}^{-1} = \begin{bmatrix} S_1^{-1} & 0 \\ 0 & S_{M,\Phi}^{-1} \end{bmatrix}, \quad S_{\Phi}^{-1} = \begin{bmatrix} S_1^{-1} & 0 \\ 0 & S_{M,\Psi}^{-1} \end{bmatrix}$$

Here, we see that both groups have the same cost of changing improvable features, as represented in the cost submatrix S_1^{-1} . However, the cost of manipulation for group Φ is higher than that of group Ψ , namely $S_{M,\Phi}^{-1} \succ S_{M,\Psi}^{-1}$.

We are now equipped to compare the costs for the two decision subjects:

$$\begin{aligned} c(x_{\phi}, \Delta(x_{\phi})) &= \frac{|w^{\top} x_{\phi}|}{\sqrt{w_{\mathbf{A}}^{\top} S_{\Phi} w_{\mathbf{A}}}} = \frac{|w^{\top} x|}{\sqrt{w_1^{\top} S_1 w_1 + w_{\mathbf{M}}^{\top} \cdot S_{M,\Phi} \cdot w_{\mathbf{M}}}} \\ c(x_{\psi}, \Delta(x_{\psi})) &= \frac{|w^{\top} x_{\psi}|}{\sqrt{w_{\mathbf{A}}^{\top} S_{\Psi} w_{\mathbf{A}}}} = \frac{|w^{\top} x|}{\sqrt{w_1^{\top} S_1 w_1 + w_{\mathbf{M}}^{\top} \cdot S_{M,\Psi} \cdot w_{\mathbf{M}}}} \end{aligned}$$

Since $S_{M,\Phi}^{-1} \succ S_{M,\Psi}^{-1}$, we have $S_{M,\Phi} \prec S_{M,\Psi}$. And since $w_{\mathbf{M}} \neq \mathbf{0}$, this implies $0 < w_{\mathbf{M}}^{\top} S_{M,\Phi} w_{\mathbf{M}} < w_{\mathbf{M}}^{\top} S_{M,\Psi} w_{\mathbf{M}}$. As a result, $c(x_{\phi}, \Delta(x_{\phi})) > c(x_{\psi}, \Delta(x_{\psi}))$ as required. \square

C.4 Additional Analysis

Proposition 7 (Correlations between Features May Reduce Cost). *For any cost matrix S^{-1} and any nontrivial classifier h , there exist indices $k, \ell \in [d_{\mathbf{A}}]$ and $\tau \in \mathbb{R}$ such that every feature vector x has lower best-response cost under the cost matrix \tilde{S}^{-1} given by*

$$\tilde{S}_{ij}^{-1} = \tilde{S}_{ji}^{-1} = \begin{cases} S_{ij}^{-1} + \tau, & \text{if } i = k, j = \ell \\ S_{ij}^{-1}, & \text{otherwise} \end{cases}$$

than under S^{-1} ; that is, $c_{\tilde{S}^{-1}}(x, \Delta(x)) < c_{S^{-1}}(x, \Delta(x))$ for all x .

Proof. Consider any cost matrix $S^{-1} \in \mathbb{R}^{d_{\mathbf{A}} \times d_{\mathbf{A}}}$ and any nontrivial classifier h (i.e. h does not assign every x the same prediction). Since S^{-1} is positive definite, so is its inverse S , and all of their diagonal entries are positive. And since h is nontrivial, it must contain a nonzero coefficient $w_i \neq 0$. Additionally, let w_j be any other coefficient.

Let $\tilde{S}^{-1} = S^{-1} + \tau(e_i e_j^{\top} + e_j e_i^{\top})$ for some constant $\tau \in \mathbb{R}$ to be set later. We claim that there exists τ such that the best-response adaptation always costs less under \tilde{S}^{-1} than S^{-1} . To do so, we compute the inverse of \tilde{S}^{-1} and invoke the closed-form cost expression given by Theorem 1.

To begin computing the inverse, note that by the Sherman-Morrison-Woodbury formula [Golub & Van Loan \(2013\)](#),

$$\tilde{S} = \left(\tilde{S}^{-1}\right)^{-1} = S - \tau S \begin{bmatrix} e_i & e_j \end{bmatrix} \left(I + \tau \begin{bmatrix} e_j^\top \\ e_i^\top \end{bmatrix} S \begin{bmatrix} e_i & e_j \end{bmatrix} \right)^{-1} \begin{bmatrix} e_j^\top \\ e_i^\top \end{bmatrix} S \quad (11)$$

$$= S - \tau S \begin{bmatrix} e_i & e_j \end{bmatrix} \left(I + \tau \begin{bmatrix} S_{ij} & S_{jj} \\ S_{ii} & S_{ij} \end{bmatrix} \right)^{-1} \begin{bmatrix} e_j^\top \\ e_i^\top \end{bmatrix} S \quad (12)$$

$$= S - \tau S \begin{bmatrix} e_i & e_j \end{bmatrix} \left[\tau \left(\frac{1}{\tau} I + \begin{bmatrix} S_{ij} & S_{jj} \\ S_{ii} & S_{ij} \end{bmatrix} \right) \right]^{-1} \begin{bmatrix} e_j^\top \\ e_i^\top \end{bmatrix} S \quad (13)$$

$$= S - \tau S \begin{bmatrix} e_i & e_j \end{bmatrix} \tau^{-1} \begin{bmatrix} \frac{1}{\tau} + S_{ij} & S_{jj} \\ S_{ii} & \frac{1}{\tau} + S_{ij} \end{bmatrix}^{-1} \begin{bmatrix} e_j^\top \\ e_i^\top \end{bmatrix} S \quad (14)$$

$$= S - S \begin{bmatrix} e_i & e_j \end{bmatrix} \underbrace{\begin{bmatrix} \frac{1}{\tau} + S_{ij} & S_{jj} \\ S_{ii} & \frac{1}{\tau} + S_{ij} \end{bmatrix}^{-1}}_T \begin{bmatrix} e_j^\top \\ e_i^\top \end{bmatrix} S \quad (15)$$

Clearly, we can ensure that T is invertible by setting τ so that $\det(T) \neq 0$. But as the following lemmas show, we can actually say much more: $\det(T)$ can be made either positive or negative, and moreover, both can be accomplished with a choice of $\tau > 0$ or $\tau < 0$. This flexibility in choosing τ will become crucial later.

First, we need the following useful fact about positive definite matrices:

Lemma 5 (Off-diagonal entries of a positive definite matrix). *If $A \in \mathbb{R}^{n \times n}$ is symmetric positive definite, then for all $i, j \in [n]$, $\sqrt{A_{ii}A_{jj}} > |A_{ij}|$.*

Proof. By positive definiteness, we have, for any nonzero $\alpha, \beta \in \mathbb{R}$,

$$(\alpha e_i + \beta e_j)^\top A (\alpha e_i + \beta e_j) = \alpha^2 A_{ii} + \beta^2 A_{jj} + 2\alpha\beta A_{ij} > 0$$

For a choice of $\alpha = -A_{ij}$ and $\beta = A_{ii}$, we have

$$A_{ij}^2 A_{ii} + A_{ii}^2 A_{jj} - 2A_{ij}^2 A_{ii} = A_{ii}(A_{ii}A_{jj} - A_{ij}^2) > 0$$

Since $A_{ii} > 0$, we must have $A_{ii}A_{jj} - A_{ij}^2 > 0$, from which the claim follows. \square

Now we can characterize the possible settings of τ and $\det(T)$:

Lemma 6 (Possible settings of τ). *There exist $\tau_{\max}, \tau_{\min} > 0$ such that the following hold:*

1. $\det(T) > 0$ for any $\tau \in \mathbb{R}$ such that $\tau_{\max} \geq |\tau| > 0$.
2. $\det(T) < 0$ for any $\tau \in \mathbb{R}$ such that $\tau_{\min} \leq |\tau|$.

Proof. To prove the first claim, note that having

$$\det(T) = \left(\frac{1}{\tau} + S_{ij} \right)^2 - S_{ii}S_{jj} > 0$$

is equivalent to

$$\left| \frac{1}{\tau} + S_{ij} \right| > \sqrt{S_{ii}S_{jj}}$$

It suffices to choose τ such that

$$\begin{aligned} \left| \frac{1}{\tau} \right| - |S_{ij}| &> \sqrt{S_{ii}S_{jj}} \\ \frac{1}{|\tau|} &> \sqrt{S_{ii}S_{jj}} + |S_{ij}| \end{aligned}$$

So any τ such that $0 < |\tau| < (\sqrt{S_{ii}S_{jj}} + |S_{ij}|)^{-1}$ results in $\det(T) > 0$. Analogously, for the second claim, a sufficient condition for $\det(T) < 0$ is that

$$\frac{1}{|\tau|} < \sqrt{S_{ii}S_{jj}} - |S_{ij}|$$

By Lemma 5, the right-hand side is positive. Hence it suffices to pick any τ such that

$$|\tau| > \left(\sqrt{S_{ii}S_{jj}} - |S_{ij}| \right)^{-1}.$$

□

With this lemma in place, we can describe the difference between the inverses of S^{-1} and \tilde{S}^{-1} . Denote this matrix by $E = S - \tilde{S}$. We show the following:

Lemma 7 (Difference between inverse cost matrices). *The k, ℓ -th entry of E has the following form:*

$$E_{k\ell} = \frac{1}{\det(T)} \left(E'_{k\ell} + \frac{1}{\tau} E''_{k\ell} \right)$$

where $E'_{k\ell}$ and $E''_{k\ell}$ do not depend on τ .

Proof. Assume that τ has been chosen so that $\det(T) \neq 0$, as Lemma 6 showed to be possible. We then have

$$T^{-1} = \frac{1}{\det(T)} \begin{bmatrix} \frac{1}{\tau} + S_{ij} & -S_{jj} \\ -S_{ii} & \frac{1}{\tau} + S_{ij} \end{bmatrix}$$

Thus continuing from equation 15, we have

$$\tilde{S} = S - \frac{1}{\det(T)} S \underbrace{\begin{bmatrix} e_i & e_j \end{bmatrix} \begin{bmatrix} \frac{1}{\tau} + S_{ij} & -S_{jj} \\ -S_{ii} & \frac{1}{\tau} + S_{ij} \end{bmatrix} \begin{bmatrix} e_j^\top \\ e_i^\top \end{bmatrix}}_V S$$

It can be verified that V is a $d_A \times d_A$ matrix whose only nonzero entries are

$$V_{ii} = -S_{jj}, \quad V_{jj} = -S_{ii}, \quad V_{ij} = V_{ji} = \frac{1}{\tau} + S_{ij}$$

Next we evaluate the $d_A \times d_A$ matrix SVS . For any $k, \ell \in [d_A]$, we have

$$\begin{aligned} (SVS)_{k\ell} &= \sum_{i'=1}^{d_A} \sum_{j'=1}^{d_A} S_{ki'} V_{i'j'} S_{j'\ell} \\ &= S_{ki} V_{ii} S_{i\ell} + S_{ki} V_{ij} S_{j\ell} + S_{kj} V_{ji} S_{i\ell} + S_{kj} V_{jj} S_{j\ell} && (V \text{ has four nonzero entries}) \\ &= V_{ii} S_{ki} S_{i\ell} + V_{jj} S_{kj} S_{j\ell} + V_{ij} (S_{ki} S_{j\ell} + S_{kj} S_{i\ell}) && (V_{ij} = V_{ji}) \\ &= -S_{jj} S_{ki} S_{i\ell} - S_{ii} S_{kj} S_{j\ell} + \left(\frac{1}{\tau} + S_{ij} \right) (S_{ki} S_{j\ell} + S_{kj} S_{i\ell}) \\ &= \underbrace{-S_{jj} S_{ki} S_{i\ell} - S_{ii} S_{kj} S_{j\ell}}_{E'_{k\ell}} + S_{ij} (S_{ki} S_{j\ell} + S_{kj} S_{i\ell}) + \frac{1}{\tau} \underbrace{(S_{ki} S_{j\ell} + S_{kj} S_{i\ell})}_{E''_{k\ell}} \end{aligned}$$

which proves the claim. □

We now compute the marginal best-response cost incurred due to the difference between the inverse cost matrices, $E = S - \tilde{S}$. We have

$$\begin{aligned}
w_A^\top E w_A &= \sum_{k=1}^{d_A} \sum_{\ell=1}^{d_A} w_k w_\ell E_{k\ell} \\
&= \frac{1}{\det(T)} \sum_{k=1}^{d_A} \sum_{\ell=1}^{d_A} w_k w_\ell \left(E'_{k\ell} + \frac{1}{\tau} E''_{k\ell} \right) && \text{(by Lemma \ref{lemma7})} \\
&= \frac{1}{\det(T)} \left[\underbrace{\sum_{k=1}^{d_A} \sum_{\ell=1}^{d_A} w_k w_\ell E'_{k\ell}}_{E'} + \frac{1}{\tau} \underbrace{\sum_{k=1}^{d_A} \sum_{\ell=1}^{d_A} w_k w_\ell E''_{k\ell}}_{E''} \right]
\end{aligned}$$

By Lemma \ref{lemma6} there exists $\tau \neq 0$ such that

$$\text{sign}(\det(T)) = -\text{sign}(E') \quad \text{and} \quad \text{sign}(\tau) = -\text{sign}(\det(T)) \cdot \text{sign}(E'')$$

Such a choice of τ results in $w_A^\top E w_A < 0$. Finally by Theorem \ref{theorem1}, we have for all x that

$$c_{\tilde{S}^{-1}}(x, \Delta_{\tilde{S}^{-1}}(x)) = \frac{|w^\top x|}{\sqrt{w_A^\top \tilde{S} w_A}} = \frac{|w^\top x|}{\sqrt{w_A^\top S w_A - w_A^\top E w_A}} < \frac{|w^\top x|}{\sqrt{w_A^\top S w_A}} = c_{S^{-1}}(x, \Delta_{S^{-1}}(x))$$

which completes the proof. \square

D Proofs and Derivations in Section \ref{section4}

D.1 Proof of Proposition \ref{prop5}

Proof. We want to show that the standard strategic risk conditioned on an unchanged true label is upper-bounded by the first term in our model designer's objective, $R_M(h)$:

$$\mathbb{E}_{x \sim \mathcal{D}} [\mathbb{1}[h(x_*) \neq y] \mid \Delta(y) = y] \leq \mathbb{E}_{x \sim \mathcal{D}} [\mathbb{1}(h(x_*^M) \neq y)]$$

We assume that the manipulating best response is more likely to result in a positive prediction than the unconstrained best response, given that the true labels do not change:

$$\mathbb{E}_{x \sim \mathcal{D}} [\mathbb{1}[h(x_*) \neq y] \mid \Delta(y) = y] \leq \mathbb{E}_{\mathcal{D}} [\mathbb{1}[h(x_*^M) \neq y] \mid \Delta_M(y) = y] \tag{16}$$

We therefore have:

$$\begin{aligned}
&\mathbb{E}_{x \sim \mathcal{D}} [\mathbb{1}(h(x_*^M) \neq y)] \\
&= \mathbb{E}_{x \sim \mathcal{D}} [\mathbb{1}(h(x_*^M) \neq y) \mid \Delta_M(y) \neq y] \cdot \Pr[\Delta_M(y) \neq y] \\
&\quad + \mathbb{E}_{x \sim \mathcal{D}} [\mathbb{1}(h(x_*^M) \neq y) \mid \Delta_M(y) = y] \cdot \Pr[\Delta_M(y) = y] \\
&= \mathbb{E}_{x \sim \mathcal{D}} [\mathbb{1}(h(x_*^M) \neq y) \mid \Delta_M(y) = y] && (\Pr[\Delta_M(y) = y] = 1) \\
&\geq \mathbb{E}_{x \sim \mathcal{D}} [\mathbb{1}(h(x_*) \neq y) \mid \Delta(y) = y] && \text{(by equation \ref{equation16})}
\end{aligned}$$

\square

D.2 Proof of Proposition 6

Proof. Let \mathcal{D}^* be the distribution induced by deploying classifier h . By the covariate shift assumption, $\Pr_{\mathcal{D}^*}(Y = y|X = x) = \Pr_{\mathcal{D}}(Y = y|X = x)$. Therefore

$$\begin{aligned} \Pr_{x \sim \mathcal{D}^*}[y(x) = +1] &= \mathbb{E}_{\mathcal{D}^*}[\mathbb{1}[y(x) = +1]] \\ &= \int \mathbb{1}[y(x) = +1] \Pr_{\mathcal{D}^*}(X = x) dx \\ &= \int \mathbb{1}[y(x) = +1] \frac{\Pr_{\mathcal{D}^*}(X = x)}{\Pr_{\mathcal{D}}(X = x)} \Pr_{\mathcal{D}}(X = x) dx \\ &= \int \mathbb{1}[y(x) = +1] \omega_h(x) \Pr_{\mathcal{D}}(X = x) dx \\ &= \mathbb{E}_{\mathcal{D}}[\omega_h(x) \mathbb{1}[y(x) = +1]] \end{aligned}$$

This implies

$$\Pr_{x \sim \mathcal{D}^*}[y(x) = +1] \geq \Pr_{x \sim \mathcal{D}}[y(x) = +1] \iff \mathbb{E}_{\mathcal{D}}[(\omega_h(x) - 1) \mathbb{1}[y(x) = +1]] \geq 0 \quad (17)$$

By similar reasoning, we have

$$\Pr_{x \sim \mathcal{D}^*}[h(x) = +1] = \mathbb{E}_{\mathcal{D}^*}[\mathbb{1}[h(x) = +1]] = \mathbb{E}_{\mathcal{D}}[\omega_h(x) \mathbb{1}[h(x) = +1]]$$

which implies

$$\Pr_{x \sim \mathcal{D}^*}[h(x) = +1] \geq \Pr_{x \sim \mathcal{D}}[h(x) = +1] \iff \mathbb{E}_{\mathcal{D}}[(\omega_h(x) - 1) \mathbb{1}[h(x) = +1]] \geq 0 \quad (18)$$

It is easy to verify that $\mathbb{E}_{x \sim \mathcal{D}}[\omega_h(x)] = 1$, and this gives us

$$\mathbb{E}_{\mathcal{D}}[(\omega_h(x) - 1) \mathbb{1}[y(x) = +1]] = \text{Cov}_{\mathcal{D}}(\omega_h(x), \mathbb{1}[y(x) = +1]) \quad (19)$$

$$\mathbb{E}_{\mathcal{D}}[(\omega_h(x) - 1) \mathbb{1}[h(x) = +1]] = \text{Cov}_{\mathcal{D}}(\omega_h(x), \mathbb{1}[h(x) = +1]) \quad (20)$$

By equation 17, equation 18, and equation 19, the condition

$$\Pr_{x \sim \mathcal{D}^*}[h(x) = +1] \geq \Pr_{x \sim \mathcal{D}}[h(x) = +1] \iff \Pr_{x \sim \mathcal{D}^*}[y(x) = +1] \geq \Pr_{x \sim \mathcal{D}}[y(x) = +1]$$

is equivalent to the condition

$$\text{Cov}_{\mathcal{D}}(\omega_h(x), \mathbb{1}[y(x) = +1]) \geq 0 \iff \text{Cov}_{\mathcal{D}}(\omega_h(x), \mathbb{1}[h(x) = +1]) \geq 0$$

□

D.3 Derivations for the model designer's objective function

Now that we have obtained a closed-form expression for both the unconstrained and improving best response from the decision subjects, we can analyze the objective function for the model designer, and the model that would be deployed at equilibrium. Recall that the objective function for the model designer is

$$\min_{w \in \mathbb{R}^{d+1}} \mathbb{E}_{x \sim \mathcal{D}}[\mathbb{1}(h(\Delta_M(x)) \neq y)] + \lambda \mathbb{E}_{x \sim \mathcal{D}}[\mathbb{1}(h(\Delta_I(x)) \neq +1)]$$

By Theorem [1](#), $h(\Delta_{\mathbf{M}}(x))$ has the closed form

$$\begin{aligned} h(\Delta_{\mathbf{M}}(x)) &= \begin{cases} +1 & \text{if } w \cdot x \geq -2\sqrt{w_{\mathbf{M}}^{\top} S_{\mathbf{M}} w_{\mathbf{M}}} \\ -1 & \text{otherwise} \end{cases} \\ &= 2 \cdot \mathbf{1} \left[w \cdot x \geq -2\sqrt{w_{\mathbf{M}}^{\top} S_{\mathbf{M}} w_{\mathbf{M}}} \right] - 1 \end{aligned}$$

and similarly,

$$h(\Delta_{\mathbf{I}}(x)) = 2 \cdot \mathbf{1} \left[w \cdot x \geq -2\sqrt{w_{\mathbf{I}}^{\top} S_{\mathbf{I}} w_{\mathbf{I}}} \right] - 1$$

The model designer's objective can then be re-written as follows:

$$\begin{aligned} &\mathbb{E}_{x \sim \mathcal{D}} [\mathbf{1}[h(\Delta_{\mathbf{M}}(x)) \neq y] + \lambda \mathbf{1}[h(\Delta_{\mathbf{I}}(x)) \neq +1]] \\ &= \mathbb{E}_{x \sim \mathcal{D}} \left[1 - \frac{1}{2}(1 + h(\Delta_{\mathbf{M}}(x)) \cdot y) + \lambda \left(1 - \frac{1}{2}(1 + h(\Delta_{\mathbf{I}}(x)) \cdot 1) \right) \right] \\ &= \mathbb{E}_{x \sim \mathcal{D}} \left[\frac{1}{2}(1 + \lambda) - \frac{1}{2}h(\Delta_{\mathbf{M}}(x)) \cdot y - \frac{\lambda}{2}h(\Delta_{\mathbf{I}}(x)) \right] \end{aligned}$$

Removing the constants, the objective function becomes:

$$\begin{aligned} &\min_w \mathbb{E}_{x \sim \mathcal{D}} [\lambda - h(\Delta_{\mathbf{M}}(x)) \cdot y - \lambda h(\Delta_{\mathbf{I}}(x))] \\ &= \min_w \mathbb{E}_{x \sim \mathcal{D}} \left[- \left(2 \cdot \mathbf{1} \left[w \cdot x \geq -2\sqrt{w_{\mathbf{M}}^{\top} S_{\mathbf{M}} w_{\mathbf{M}}} \right] - 1 \right) \cdot y(x) - 2\lambda \cdot \mathbf{1} \left[w \cdot x \geq -2\sqrt{w_{\mathbf{I}}^{\top} S_{\mathbf{I}} w_{\mathbf{I}}} \right] \right] \end{aligned}$$

Re-organizing the above equations, we can turn the model designer's *constrained* optimization problem in equation [6](#) into the following *unconstrained* problem:

$$\min_{w \in \mathbb{R}^d} \mathbb{E}_{x \sim \mathcal{D}} \left[- \left(2 \cdot \mathbf{1} \left[w^{\top} x \geq -2\sqrt{\Omega_{\mathbf{M}}} \right] - 1 \right) \cdot y - 2\lambda \cdot \mathbf{1} \left[w^{\top} x \geq -2\sqrt{\Omega_{\mathbf{I}}} \right] \right] \quad (21)$$

The optimization problem in equation [21](#) is intractable since both the objective and the constraints are non-convex. To overcome this difficulty, we train our classifier by replacing the 0-1 loss function with a convex surrogate loss $\sigma(x) = \log\left(\frac{1}{1+e^{-x}}\right)$. This results in the following ERM problem:

$$\min_{w \in \mathbb{R}^d} \frac{1}{n} \sum_{i=1}^n \left[-\sigma\left(y_i \cdot (w^{\top} x_i + 2\sqrt{\Omega_{\mathbf{M}}})\right) - \lambda \cdot \sigma(w^{\top} x_i + 2\sqrt{\Omega_{\mathbf{I}}}) \right] \quad (22)$$

Conditionally Actionable Features. In practice, individuals can often only change some features in either a positive or negative direction, but not both. However, modeling this restriction on the decision subject's side precludes a closed-form solution. Instead, we strongly disincentivize such moves in the model designer's objective function. The idea is that if the model designer is punished for encouraging an illegal action, the announced classifier will not incentivize such moves from decision subjects. The result is that decision subjects encounter an *implicit* direction constraint on the relevant variables. To that end, we construct a vector $\text{dir} \in \{-1, 0, +1\}^d$ where dir_i represents the prohibited direction of change for the corresponding feature x_i ; that is, $\text{dir}_i = +1$ if x_i should not be allowed to increase, -1 if it should not decrease, and 0 if there are no direction constraints. We then append the following penalty term to the model designer's objective in Eq. [6](#):

$$-\eta \cdot \sum_{i=1}^d \max(\text{dir}_i \cdot (\Delta(x) - x)_i, 0) \quad (23)$$

where $\eta > 0$ is a hyperparameter representing the weight given to this penalty term. Eq. [23](#) penalizes the weights of partially actionable features so that decision subjects would prefer to move towards a certain direction.

E Additional Related Work

Strategic Classification. There has been extensive research on strategic behavior in classification [Hardt et al. \(2016a\)](#); [Cai et al. \(2015\)](#); [Chen et al. \(2018\)](#); [Dong et al. \(2018\)](#); [Dekel et al. \(2010\)](#); [Chen et al. \(2020\)](#). [Hardt et al. \(2016a\)](#) was the first to formalize strategic behavior in classification based on a sequential two-player game (i.e. a Stackelberg game) between decision subjects and classifiers. Since then, other similar Stackelberg formulations have been studied [Balcan et al. \(2015\)](#). [Dong et al. \(2018\)](#) considers the setting in which decision subjects arrive in an online fashion and the learner lacks full knowledge of decision subjects' utility functions. More recently, [Chen et al. \(2020\)](#) proposes a learning algorithm with non-smooth utility and loss functions that adaptively partitions the learner's action space according to the decision subject's best responses.

Recourse. The concept of *recourse* in machine learning was first introduced in [Ustun et al. \(2019\)](#). There, an integer programming solution was developed to offer actionable recourse from a linear classifier. Our work builds on theirs by considering strategic actions from decision subjects, as well as by aiming to incentivize honest improvement. [Venkatasubramanian & Alfano \(2020\)](#) discusses a more adequate conceptualization and operationalization of recourse. [Karimi et al. \(2020a\)](#) provides a thorough survey of algorithmic recourse in terms of its definitions, formulations, solutions, and prospects. Inspired by the concept of recourse, [Dean et al. \(2020\)](#) develops a reachability problem to capture the ability of models to accommodate arbitrary changes in the interests of individuals in recommender systems. [Bellamy et al. \(2018\)](#) builds toolkits for actionable recourse analysis. Furthermore, [Gupta et al. \(2019\)](#) studies how to mitigate disparities in recourse across populations.

Causal Modeling of Features. A flurry of recent papers have demonstrated the importance of understanding causal factors for achieving fairness in machine learning [Wang et al. \(2019\)](#); [Bhatt et al. \(2020\)](#); [Bechavod et al. \(2020\)](#); [Miller et al. \(2020\)](#); [Shavit et al. \(2020\)](#). [Miller et al. \(2020\)](#) studies distinctions between gaming and improvement from a causal perspective. [Shavit et al. \(2020\)](#) provides efficient algorithms for simultaneously minimizing predictive risk and incentivizing decision subjects to improve their outcomes in a linear setting. In addition, [Karimi et al. \(2020b\)](#) develops methods for discovering recourse-achieving actions with high probability given limited causal knowledge. In contrast to these works, we explicitly separate improvable features from manipulated features when maximizing decision subjects' payoffs.

Incentive Design. Like our work, [Kleinberg & Raghavan \(2020\)](#) discusses how to incentivize decision subjects to improve a certain subset of features. Next, [Haghtalab et al. \(2020\)](#) shows that an appropriate projection is an optimal linear mechanism for strategic classification, as well as an *approximate* linear threshold mechanism. Our work complements theirs by providing appropriate linear classifiers that balance accuracy and improvement. [Liu et al. \(2020\)](#) considers the equilibria of a dynamic decision-making process in which individuals from different demographic groups invest rationally, and compares the impact of two interventions: decoupling the decision rule by group and subsidizing the cost of investment.

Algorithmic Fairness in Machine Learning. Our work contributes to the broad study of algorithmic fairness in machine learning. Most common notions of group fairness include disparate impact [Feldman et al. \(2015\)](#), demographic parity [Agarwal et al. \(2018\)](#), disparate mistreatment [Zafar et al. \(2019\)](#), equality of opportunity [Hardt et al. \(2016b\)](#) and calibration [Chouldechova \(2017\)](#). Among them, disparities in the recourse fraction can be viewed as equality of false positive rate (FPR) in the strategic classification setting. Disparities in costs and flipsets are also relevant to counterfactual fairness [Kusner et al. \(2017\)](#) and individual fairness [Dwork et al. \(2012\)](#). Similar to our work, [von Kügelgen et al. \(2020\)](#) also consider the intervention cost of recourse in flipping the prediction across subgroups, investigating the fairness of recourse from a causal perspective.

E.1 Agent’s Best Response with Partially Actionable Features

Let feature i represents those features that should only be non-increasing, and feature j represents those features that should only be non-decreasing. Then the constraint can be represented as:

$$\begin{aligned} y_i \leq 0 &\Leftrightarrow e_i^\top y \leq 0 \\ y_j \geq 0 &\Leftrightarrow e_j^\top y \geq 0 \end{aligned}$$

Assume that there are n_- features that can only be changed negatively, and there are n_+ features that can only be changed increasingly. We can further combine those new constraints into a matrix form like $Ey \leq 0$. The other constraint can be re-written as:

$$w^\top y - b' \geq 0 \Leftrightarrow -w^\top y \leq -b',$$

therefore the optimization problem can be rewritten as:

$$\begin{aligned} \min \quad & \frac{1}{2} y^\top Q y \\ \text{s.t.} \quad & \underbrace{\begin{bmatrix} E \\ -w^\top \end{bmatrix}}_A y \leq \underbrace{\begin{bmatrix} 0 \\ -b' \end{bmatrix}}_b \end{aligned}$$

where A is of the form:

$$A = \begin{bmatrix} I_{n_- \times n_-} & 0 & 0 \\ 0 & -I_{n_+ \times n_+} & 0 \\ \ddots & -w^\top & \ddots \end{bmatrix}_{(n_+ + n_- + 1) \times n}$$

F Additional Experimental Details and Results

In this section, we provide additional experimental results. In particular, we provide the full results with mean and standard deviation in Table 4.

F.1 Basic information of each dataset

Table 3: Basic information of each dataset.

Dataset	Size	Dimension	Prediction Task
credit	20,000	16	To predict if a person can repay their credit card loan.
adult	48,842	14	To predict whether income exceeds 50K/yr based on census data.
german	1,000	26	To predict whether a person is good or bad credit risk.
spam	4601	57	To predict if an email is a spam or not.

F.2 Specific Cost Matrix

We specify the cost matrix as follows:

$$S_{ij}^{-1} = \begin{cases} 1, & \text{if } i = j \text{ and } i \in \text{I} \\ 0.2, & \text{if } i = j \text{ and } j \in \text{M} \\ 1, & \text{if the cost of changing features } i \\ & \text{and } j \text{ are } \textit{negatively} \text{ correlated} \\ -1, & \text{if the cost of changing features } i \\ & \text{and } j \text{ are } \textit{positively} \text{ correlated} \\ 0, & \text{otherwise} \end{cases}$$

Table 4: Performance metrics for all methods over 4 real data sets with non-diagonal cost matrix. We report the mean and standard deviation for 5-fold cross validation. The constructive adaptation (CA) consistently achieves a high accuracy at deployment while providing the highest improvement rates across all four datasets.

Dataset	Metrics	METHODS			
		ST	DF	MP	CA
CREDIT	<i>test error</i>	29.52 ± 0.37	29.66 ± 0.40	29.65 ± 0.41	29.60 ± 0.44
	<i>deploy error</i>	31.25 ± 0.56	29.66 ± 0.40	29.41 ± 0.32	29.49 ± 0.38
	<i>improvement rate</i>	46.35 ± 3.81	44.71 ± 4.75	36.76 ± 0.53	48.27 ± 5.50
ADULT	<i>test error</i>	23.05 ± 0.47	33.55 ± 0.73	24.94 ± 0.52	27.22 ± 0.65
	<i>deploy error</i>	38.64 ± 4.46	33.55 ± 0.73	26.85 ± 0.59	29.34 ± 0.45
	<i>improvement rate</i>	30.92 ± 3.31	60.63 ± 29.40	36.70 ± 1.62	63.79 ± 7.80
GERMAN	<i>test error</i>	30.85 ± 0.82	36.10 ± 1.97	33.25 ± 1.44	34.70 ± 2.15
	<i>deploy error</i>	33.40 ± 1.78	36.10 ± 1.97	34.60 ± 1.94	34.25 ± 1.78
	<i>improvement rate</i>	41.20 ± 5.77	42.10 ± 9.07	33.50 ± 2.53	56.10 ± 6.40
SPAMBASE	<i>test error</i>	7.11 ± 0.52	10.18 ± 0.45	11.52 ± 0.12	14.37 ± 0.24
	<i>deploy error</i>	22.40 ± 3.14	10.18 ± 0.45	12.92 ± 0.58	14.70 ± 0.36
	<i>improvement rate</i>	40.04 ± 13.06	32.46 ± 14.63	26.42 ± 4.80	43.98 ± 6.18

We use the `credit` dataset as a demonstration of how we specify the non-diagonal element in the cost matrix. For two feature variables that have a positive correlation, e.g., *CheckingAccountBalance* and *SavingsAccountBalance*, we assign -1 to the corresponding elements in the cost matrix S . For two feature variables that have a negative correlation, e.g., *CheckingAccountBalance* and *MissedPayments*, we assign $+1$ to the corresponding elements in the cost matrix.

F.3 Computing Infrastructure

We conducted all experiments on a 3 GHz 6-Core Intel Core i5 CPU. All our methods have relatively modest computational cost and can be trained within a few minutes.

F.4 Results for non-diagonal cost matrix

In real life, the specification of the cost matrix might require examining the causal correlations among different features. We consider a non-diagonal cost matrix setup based on common knowledge and describe the rationale as below. For two feature variables that have a positive correlation, e.g., *CheckingAccountBalance* and *SavingsAccountBalance*, we assign -1 to the corresponding elements in the cost matrix. For two feature variables that have a negative correlation, e.g., *CheckingAccountBalance* and *MissedPayments*, we assign $+1$ to the corresponding elements in the cost matrix. We also note that the non-diagonal cost matrix must be invertible under our assumption on the cost of modifying features. We provide more detailed results for each dataset in Table 4, which shows the means and standard deviations of different metrics. Compared to the empirical results of using a diagonal matrix, we achieve similar results with respect to the three evaluation criteria across all four methods.

F.5 Additional Experimental Results for Non-Linear models

We also work with a three-layer neural network to validate the effectiveness of the oracle best response in Algorithm 1. We note that the LIME program needs to learn a local linear model for each instance, which is very time-consuming. Therefore, we downsample only 10% of data examples from the `credit` dataset. We follow the same setting as the linear classifier experiments. We compare our method with the static classifier in Table 5. We find out for this non-linear model setting, our approach has a higher improvement rate while preventing manipulations with the deploy error 27.72% vs. 35.64%.

Table 5: Performance metrics for non-linear models.

Metrics	METHODS	
	ST	CA
<i>test error</i>	30.72%	30.01%
<i>deploy error</i>	35.64%	27.72%
<i>improvement rate</i>	0.99%	2.97%

F.6 Flipsets

We also construct flipsets for individuals in the `german` dataset using the closed-form solution Eq. (3) under our trained classifier. The individual characterized as a “bad consumer” (-1) is supposed to decrease their missed payments in order to flip their outcome of the classifier with respect to a non-diagonal cost matrix. In contrast, even though the individual improves their loan rate or liable individuals, the baseline classifier will still reject them. We also provide flipsets for conditionally actionable features on the `credit` dataset in Table 7. The individual will undesirably reduce their education level when the classifier is unaware of the partially actionable features. In contrast, the individual decreases their total overdue months instead when the direction penalty is imposed during training.

Table 6: Flipset for a person denied credit by ManipulatedProof on the `german` dataset. The red up arrows \uparrow represent increasing the values of features, while the red down arrows \downarrow represent decreasing.

Feature	Type	Original	LightTouch	ManipulatedProof
<i>LoanRateAsPercentOfIncome</i>	I	3	3	2 \downarrow
<i>NumberOfOtherLoansAtBank</i>	I	1	1	1
<i>NumberOfLiableIndividuals</i>	I	1	0 \downarrow	2 \uparrow
<i>CheckingAccountBalance</i> ≥ 0	I	0	0	0
<i>CheckingAccountBalance</i> ≥ 200	I	0	0	0
<i>SavingsAccountBalance</i> ≥ 100	I	0	0	0
<i>SavingsAccountBalance</i> ≥ 500	I	0	0	0
<i>MissedPayments</i>	I	1	0 \downarrow	1
<i>NoCurrentLoan</i>	I	0	0	0
<i>CriticalAccountOrLoansElsewhere</i>	I	0	0	0
<i>OtherLoansAtBank</i>	I	0	0	0
<i>OtherLoansAtStore</i>	I	0	0	0
<i>HasCoapplicant</i>	I	0	0	0
<i>HasGuarantor</i>	I	0	0	0
<i>Unemployed</i>	I	0	0	0
<i>LoanDuration</i>	M	48	47 \downarrow	47 \downarrow
<i>PurposeOfLoan</i>	M	0	0	0
<i>LoanAmount</i>	M	4308	4307 \downarrow	4307 \downarrow
<i>HasTelephone</i>	M	0	0	0
<i>Gender</i>	U	0	0	0
<i>ForeignWorker</i>	U	0	0	0
<i>Single</i>	U	0	0	0
<i>Age</i>	U	24	24	24
<i>YearsAtCurrentHome</i>	U	4	4	4
<i>OwnsHouse</i>	U	0	0	0
<i>RentsHouse</i>	U	1	1	1
<i>YearsAtCurrentJob</i> ≤ 1	U	1	1	1
<i>YearsAtCurrentJob</i> ≥ 4	U	0	0	0
<i>JobClassIsSkilled</i>	U	1	1	1
<i>GoodConsumer</i>	-	-1	+1 \uparrow	-1

Table 7: Flipset for an individual on Credit dataset with partially actionable features. The red up arrows \uparrow represent any increasing values, while the red down arrows \downarrow represent any decreasing values.

Feature	Type	dir	Original	$\eta = 0$	$\eta = 100$
<i>EducationLevel</i>	I	+1	3	2 \downarrow	3
<i>TotalOverdueCounts</i>	I	0	1	1	1
<i>TotalMonthsOverdue</i>	I	0	1	1	0 \downarrow
<i>MaxBillAmountOverLast6Months</i>	M	0	0	0	0
<i>MaxPaymentAmountOverLast6Months</i>	M	0	0	0	0
<i>MonthsWithZeroBalanceOverLast6Months</i>	M	0	0	0	0
<i>MonthsWithLowSpendingOverLast6Months</i>	M	0	6	5 \downarrow	6
<i>MonthsWithHighSpendingOverLast6Months</i>	M	0	0	0	0
<i>MostRecentBillAmount</i>	M	0	0	0	0
<i>MostRecentPaymentAmount</i>	M	0	0	0	0
<i>Married</i>	U	0	1	1	1
<i>Single</i>	U	0	0	0	0
<i>Age \leq 25</i>	U	0	0	0	0
<i>25 \leq Age \leq 40</i>	U	0	0	0	0
<i>40 \leq Age $<$ 60</i>	U	0	0	0	0
<i>Age \geq 60</i>	U	0	1	1	1
<i>HistoryOfOverduePayments</i>	U	0	1	1	1
<i>NoDefaultNextMonth</i>	-	-	-1	+1 \uparrow	+1 \uparrow