

1 The Role of Randomness and Noise in Strategic 2 Classification

3 **Mark Braverman**

4 Department of Computer Science, Princeton University, USA

5 mbraverm@cs.princeton.edu

6 **Sumegha Garg**

7 Department of Computer Science, Princeton University, USA

8 sumeghag@cs.princeton.edu

9 — Abstract —

10 We investigate the problem of designing optimal classifiers in the “strategic classification” setting,
11 where the classification is part of a game in which players can modify their features to attain a
12 favorable classification outcome (while incurring some cost). Previously, the problem has been
13 considered from a learning-theoretic perspective and from the algorithmic fairness perspective.

14 Our main contributions include

- 15 ■ Showing that if the objective is to maximize the efficiency of the classification process (defined
16 as the accuracy of the outcome minus the sunk cost of the qualified players manipulating their
17 features to gain a better outcome), then using randomized classifiers (that is, ones where the
18 probability of a given feature vector to be accepted by the classifier is strictly between 0 and 1)
19 is necessary.
- 20 ■ Showing that in many natural cases, the imposed optimal solution (in terms of efficiency) has
21 the structure where players never change their feature vectors (and the randomized classifier is
22 structured in a way, such that the gain in the probability of being classified as a ‘1’ does not
23 justify the expense of changing one’s features).
- 24 ■ Observing that the randomized classification is not a *stable* best-response from the classifier’s
25 viewpoint, and that the classifier doesn’t benefit from randomized classifiers without creating
26 instability in the system.
- 27 ■ Showing that in some cases, a *noisier signal* leads to better equilibria outcomes — improving
28 both accuracy and fairness when more than one subpopulation with different feature adjustment
29 costs are involved. This is particularly interesting from a policy perspective, since it is hard to
30 force institutions to stick to a particular randomized classification strategy (especially in a context
31 of a market with multiple classifiers), but it is possible to alter the information environment to
32 make the feature signals inherently noisier.

33 **2012 ACM Subject Classification** Theory of computation → Algorithmic game theory and mechanism
34 design; Theory of computation → Machine learning theory

35 **Keywords and phrases** Strategic classification, noisy features, randomized classification, fairness

36 **Digital Object Identifier** 10.4230/LIPIcs.FORC.2020.9

37 **Funding** *Mark Braverman*: Research supported in part by the NSF Alan T. Waterman Award,
38 Grant No. 1933331, a Packard Fellowship in Science and Engineering, and the Simons Collaboration
39 on Algorithms and Geometry. Any opinions, findings, and conclusions or recommendations expressed
40 in this publication are those of the author and do not necessarily reflect the views of the National
41 Science Foundation.

42 **1 Introduction**

43 Machine learning algorithms are increasingly being used to make decisions about the individuals in various areas such as university admissions, employment, health, etc. As the individuals gain information about the algorithms being used, they have an incentive to



© M. Braverman and S. Garg;

licensed under Creative Commons License CC-BY

1st Symposium on Foundations of Responsible Computing (FORC 2020).

Editor: Aaron Roth; Article No. 9; pp. 9:1–9:20



Leibniz International Proceedings in Informatics

LIPIcs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

9:2 The Role of Randomness and Noise in Strategic Classification

46 adapt their data so as to be classified desirably. For example, if a student is aware that
47 a university heavily weighs SAT score in their admission process, she will be motivated
48 to achieve a higher SAT score either through extensive test preparation or multiple tries.
49 Such efforts by the students might not change their probability of being successful at the
50 university, but are enough to fool the admissions' process. Therefore, under such "strategic
51 manipulation" of one's data, the predictive power of the decisions are bound to decrease.
52 One way to prevent such manipulation is by keeping the classification algorithms a secret,
53 but this is not a practical solution to the problem, as some information is bound to leak
54 over time and the transparency of these algorithms is a growing social concern. Thus, this
55 motivates the study of algorithms that are optimal under "strategic manipulation". The
56 problem of gaming in the context of classification algorithms is a well known problem and is
57 increasingly gaining researchers' attention, for example, [8, 1, 9, 16, 4].

58 [2] and [8] modeled strategic classification as a Stackelberg competition– the algorithm
59 (Jury) goes first and publishes the classifier, and then the individuals get to transform their
60 data, after knowing the classifier, incurring certain costs to manipulate. The individuals
61 would manipulate their features as long as the cost to manipulate is less than the advantage
62 gained in getting the desirable classification. We assume that such manipulations don't
63 change the actual qualifications of an individual. A natural question is: what classifier
64 achieves optimal classification accuracy under the Stackelberg competition? These papers
65 considered the task of strategic classification when the published classifier is deterministic. We
66 study the role of randomness (and addition of noise to the features) in strategic classification
67 and define the Stackelberg equilibrium for probabilistic classifiers, that assigns a real number
68 in $[0, 1]$, to each individual and a classification outcome o , representing the probability of
69 being classified as o .

70 As higher SAT scores are preferred by a university, the students would put an effort in
71 increasing their SAT score, thereby, forcing the university to raise the score bar to optimize
72 its accuracy (under the Stackelberg equilibrium). Due to this increased bar of acceptance,
73 even the students who were above the true cutoff would have to put an extra effort to
74 achieve a SAT score above this raised bar. And this effort is entirely the result of gaming
75 in the classification system. We define the *cost of strategy* for a published classifier to be
76 the total extra effort, it induced, amongst the qualified individuals of the population. Then,
77 we define the *efficiency* of a published classifier to be its classification accuracy minus the
78 cost of strategy under the Stackelberg equilibrium. A natural question here is: what classifier
79 achieves the optimal efficiency? The efficiency of a published classifier represents the total
80 impact of the classifier on all the agents in the Stackelberg equilibrium.

81 In normal classification problems it is never a good idea to use randomness, since one
82 should always adhere to the best/utility maximizing action based on the prediction. Just as
83 in games, randomness may lead to better solution in strategic classification, the paper aims
84 to start understanding tradeoffs between efficiency losses due to randomness and efficiency
85 gains through better equilibria induced by the randomized classifier.

86 Gaming in classification adds to the plethora of fairness concerns associated with classi-
87 fication algorithms, when the costs of manipulation are different across subpopulations. For
88 example, a high weightage of SAT scores (for university admissions) favors the subgroups of
89 the society that have the resources to enroll in test preparation or attempt the test multiple
90 times. Further, varying costs across the subpopulations can lead to varied efforts put by
91 identically qualified individuals, belonging to different subpopulations, to achieve the same
92 outcome. [16] and [9] study the disparate effects of strategic classification on subpopulations
93 (we will discuss these papers more in the related work section). [9] observes that a single

94 classifier might have different classification errors on subpopulations due to the varying cost
 95 of manipulations. We also study the effect of strategic manipulation on the classification
 96 errors across subpopulations and how randomized classifiers or noisy features may reduce
 97 the disparate effects.

98 Strategic classification is a well known problem and there has been research in many other
 99 aspects of strategic classification, for example, learning the optimal classifier efficiently when
 100 the samples might also be strategic [8, 4], mechanism design under strategic manipulation
 101 [3, 5, 12], and studying the manipulation costs that actually change the inherent qualifications
 102 [14, 15]. The focus of this paper is theoretically demonstrating the role of randomness and
 103 noise in the strategic setting.

104 1.1 Our contributions

105 Above, we talked about how strategic manipulation can deteriorate the classification accuracy
 106 and lead to unfair classification. We investigate the different scenarios of the classification
 107 task that help in regaining the lost accuracy and fairness guarantees. Our entire work is
 108 based on *one-dimensional feature space*.

109 1.1.1 Randomized classifiers

110 Firstly, we formulate the strategic classification task, when the published classifier is ran-
 111 domized. Instead of publishing a single binary classifier (for 2 classification outcomes, 0 and
 112 1), the Jury publishes a distribution of classifiers and promises to pick the final classifier
 113 from that distribution. Another interpretation is that the Jury assigns a value in $[0, 1]$ to
 114 each feature value, which represents the probability of an individual with this feature being
 115 classified as 1. The individuals manipulate their features, after knowing the set of classifiers
 116 but not the final classifier, incurring certain costs according to the *cost function*.

117 Not surprisingly, we show through examples that a probabilistic classifier can achieve
 118 strictly higher expected accuracy and efficiency than any binary classifier under strategic
 119 setting. Note that, without any strategic manipulation, a randomized classifier has no
 120 advantage over deterministic classifiers in terms of classification accuracy. The intuition is as
 121 follows: using randomness, the Jury can discourage the individuals from manipulating their
 122 features by making the advantage gained by any such a manipulation small enough.

123 For *simple* cost functions, we then characterize the randomized classifier that achieves
 124 optimal efficiency. We prove that such a classifier sets the probabilities (of being classified as
 125 1) such that none of the individuals have an incentive to manipulate their feature. Given
 126 two features x and x' in the feature space, let $c(x, x')$ denote the cost of manipulating one's
 127 feature from x to x' . Informally, we say a cost function c is *simple* when all the costs are
 128 non-negative, the cost to manipulate to a "less" qualified feature is 0, and the costs are
 129 sub-additive, that is, manipulating your feature x directly to x'' is at least easier than first
 130 manipulating it to x' and then to x'' . The characterization theorem, stated informally, is as
 131 follows:

132 ▶ **Theorem 1** (Informal statement of Theorem 3). *For simple cost functions, the most efficient
 133 randomized classifier is such that the best response of all the individuals is to reveal their
 134 true features.*

135 This characterization, in addition to being mathematically clean, allows us to infer the
 136 following: let A and B be two subpopulations (identical in terms of qualifications) such
 137 that the costs to manipulation are *higher* for individuals in A than in B , then the optimal
 138 efficiency obtained for the subpopulation A is greater than that in B .

139 **1.1.2 Obstacles to using a randomized classifier**

140 Till now, we have argued the benefits of using a probabilistic classifier. However, the degree to
 141 which it is possible to use or commit to a randomized strategy varies depending on the setting.
 142 There are two main drivers impeding the implementation of the most efficient Stackelberg
 143 equilibrium. Firstly, in many real-life classification settings, it might be unacceptable to
 144 use a probabilistic classifier, for example, due to legal restrictions (applicants with identical
 145 features must obtain identical outcomes). Secondly, for the more complicated scenario with
 146 multiple classifiers (such as college admissions), the effect of each Jury on the overall market
 147 is small, hence, diminishing the incentive to stick to a randomized strategy ‘for the benefit
 148 of the market as a whole’. Informally, the best response of a single Jury, when the other
 149 classifiers commit to using a randomized classifier, is not a randomized classifier. And even
 150 if we got the Juries to commit to randomization, the final probabilities of classification
 151 depends on the number of classifiers (k) and hence, the implementation of the most efficient
 152 randomized classifier needs coordination between the multiple classifiers. Analyzing the
 153 equilibria for multiple classifiers is beyond the scope of this paper but we illustrate the
 154 instability of randomized classifier as follows. We show that unless Jury is able to commit
 155 to the published randomized classifier, such a classifier is not a stable solution to strategic
 156 classification. As mentioned above, randomization helps because of the following observation:
 157 if the difference between the probabilities, of being classified as 1 at *adjacent* features is
 158 small, the individuals have no incentive to manipulate their features. But, once the Jury
 159 knows that no one changed their feature, her best response, then, is to use the classifier that
 160 achieves best accuracy given the *true* features.

161 Formally, we show (Theorem 5) that for any published randomized classifier that achieves
 162 strictly higher accuracy compared to any deterministic classifier under Stackelberg equilibrium,
 163 Jury has an opportunity to improve its utility and get strictly better accuracy using a classifier
 164 different from the published.

165 The shortcomings of a randomized classifier can be redeemed by addition of noise to the
 166 features.

167 **1.1.3 Addition of noise to the features**

168 This brings us to our second scenario that uses noisy features for classification. Every
 169 individual has an associated private signal that identifies their qualification. The Jury sees a
 170 feature that is a noisy representation of this private signal. The individuals, after incurring
 171 certain cost, can effectively manipulate their private signal such that the features are a
 172 noisy representation of this updated private signal. Again, the assumption is that such a
 173 manipulation didn’t change the true qualifications of an individual. We show, through an
 174 example of a cost function and a noise distribution, that in the strategic setting, using a
 175 deterministic classifier, the Jury achieves better accuracy when the features are noisy than
 176 any deterministic classifier in the noiseless case, that is, when Jury gets to see the private
 177 signal. This is counter-intuitive at first glance because under no strategic manipulation, noise
 178 can only decrease Jury’s accuracy.

179 We also show examples where noisy features can help in achieving fairer outcomes across
 180 subpopulations. Let A and B be two subpopulations *identical* in qualifications but having
 181 different (but not extremely different) costs of manipulation (and $|A| \leq |B|$; A is a minority).
 182 We show, through an example, that no matter whether the minority has higher or lower
 183 costs of manipulation than the majority, it is at a disadvantage when Jury publishes a single
 184 deterministic classifier to optimize its overall accuracy (noiseless strategic setting). Here, by

185 disadvantage, we mean that the minority has lower classification accuracy than the majority.
 186 Next, we show that the addition of appropriate noise to the private signals, in the same
 187 example, can ensure that Jury's best response classifier is fair across subpopulations. This is
 188 not that surprising as making the features completely noisy also lead to same outcomes for the
 189 subpopulations. However, such an addition of noise can also sometimes increase Jury's overall
 190 accuracy (improving both accuracy and fairness). We consider the case where the Jury would
 191 publish a single classifier for both the subpopulations (for e.g., either because A is a protected
 192 group and the Jury is not allowed to discriminate based on the subgroup membership or
 193 because the Jury has not yet identified these subpopulations and the differences in their cost
 194 functions). Informally, our results, can be stated as follows:

195 **► Theorem 2** (Informal statement of Theorems 6,7,8). *Let A and B be two subpopulations that
 196 are identical in qualifications. Let $c_A \neq c_B$ be the cost functions for subpopulations A and B
 197 respectively. In Case 1, Jury gets to see the private signals and publishes a single deterministic
 198 classifier that achieves optimal overall accuracy (sum over the two subpopulations) under
 199 the Stackelberg equilibrium (for the cost functions c_A and c_B). In Case 2, the features are
 200 noisy representations of the private signal; Jury publishes a single deterministic classifier
 201 that achieves optimal overall accuracy under the Stackelberg equilibrium (knowing that the
 202 features are noisy). There exists an instantiation of the "identical qualifications" such that*
 203 1. *If $|A| < |B|$, that is, A is a minority, for a wide set of costs functions c_A, c_B , A is always
 204 at a disadvantage when in Case 1.*
 205 2. *There exists a setting of the "noise" (η) for each of the above cost functions, such that,
 206 Jury's best response in Case 2, is always fair, that is, achieves equal classification accuracy
 207 on the subpopulations.*
 208 3. *There exists cost functions c_A, c_B from this wide set of cost functions, and corresponding
 209 noise η , such that Jury's accuracy in Case 2 is strictly better than in Case 1.*

210 This result has potentially interesting policy implications, since it is easier, both practically
 211 and legally, to commit to using noisier signals (for example by restricting the types of
 212 information available to the Jury) than to commit to disregarding pertinent information
 213 ex-post (as in randomized classification). Therefore, future mechanism design efforts involving
 214 strategic classification should carefully consider the mechanisms of information disclosure to
 215 the Jury.

216 1.2 Related Work

217 [8, 2] initiated the study of strategic classification through the lens of Stackelberg competition.
 218 [9, 16, 10] study the effects of strategic classification on different subpopulations and how it
 219 can exacerbate the social inequity in the world. [9] also made the observation that a single
 220 classifier would have varying classification accuracies across subpopulations with different
 221 costs of manipulation. [16] defined a concept called "social burden" of a classifier to be the
 222 sum of the minimum effort any qualified individual has to put in to be classified as 1. Thus,
 223 the subpopulations with higher costs of manipulation would have worse social burden and
 224 might be at a disadvantage. In such situations, intuitively, one would think that subsidizing
 225 the costs for the disadvantaged population might help. [9] showed that cost subsidy for
 226 disadvantaged individuals can sometimes lead to worse outcomes for the disadvantaged group.

227 In the present paper, we observe that the addition of noise, counter-intuitively, can help
 228 Jury's accuracy as well as serve the fairness concerns. There are many examples in game
 229 theory where loss of information helps an individual in strategic setting, for example, [6].
 230 [11, 10] also studies the role of hiding information to serve fairness. [7] has a brief discussion

231 at the end of the paper on making manipulated data more informative through addition of
 232 noise to the features (this was put online a couple of months after the first version of our
 233 paper was made online).

234 Another work related to Theorem 3 of the present paper is [13], which studies the scope
 235 of truthful mechanisms when the agents incur certain costs for misreporting their true type.
 236 In particular, the paper gives conditions, on the misreporting costs, that allow the revelation
 237 principle to hold, that is, any mechanism can be implemented by a truthful mechanism,
 238 where all the agents reveal their true types. The main difference between [13] and our paper
 239 is that the former allows the use of monetary transfers to the agents to develop truthful
 240 mechanisms and such transfers don't impact the objective value of the mechanism.

241 1.3 Organization

242 We formalize the model used for strategic classification in Section 2. In Section 3, we show
 243 how randomness helps in achieving better accuracy and efficiency. We also characterize
 244 the classifiers that achieve optimal efficiency for *simple* cost functions. In Section 4, we
 245 investigate the stability of randomized classifiers. In Section 5, we investigate the role of
 246 noisy features in strategic classification.

247 2 Preliminaries

In this paper, we concern ourselves with classification based on a one-dimensional feature space \mathcal{X} . In many of the examples, our feature space $\mathcal{X} \subseteq \mathbb{R}$ is discrete, hence, we use sum (\sum) in many of the definitions, but, these definitions are well-defined when \mathcal{X} is taken to be continuous (for e.g., \mathbb{R}) by replacing sum (\sum) with integrals (\int) and probability distributions with probability density functions. We use the notation $\mathcal{N}(z, \sigma)$ to denote the gaussian distribution with mean z and standard deviation σ . We say a function $f : \mathcal{X} \rightarrow \{0, 1\}$ is a threshold function (classifier) with threshold τ if

$$f(x) = \begin{cases} 1 & \text{if } x \geq \tau \\ 0 & \text{otherwise} \end{cases}$$

248 We also use $1_{x \geq \tau}$ to denote a threshold function (classifier) with threshold τ . Sometimes, we
 249 will use $1_{x > \tau}$ that classifies x as 1 if and only if $x > \tau$.

250 2.1 The Model

251 Let \mathcal{X} be the set of features. Let $\pi : \mathcal{X} \rightarrow [0, 1]$ be the probability distribution over the
 252 feature set realized by the individuals. Let $h : \mathcal{X} \rightarrow [0, 1]$ be the true probability of an
 253 individual being qualified (1) given the feature. We also refer to it as the true qualification
 254 function. Let $c(x, x')$ be the cost incurred by an individual to manipulate their feature from x
 255 to x' (We also use words, change and move, to refer to this manipulation). The classification
 256 is modeled as a sequential game where a Jury publishes a classifier (possibly probabilistic)
 257 $f : \mathcal{X} \rightarrow [0, 1]$ and contestants (individuals) can change their features (after seeing f) as long
 258 as they are ready to incur the cost of change. The previous papers in the area considered the
 259 task of strategic classification when the published classifier is deterministic binary classifier.
 260 Here, we formalize the Stackelberg prediction game for probabilistic classifiers.

261 Given f , we define the best response of a contestant with feature x^1 , as follows

$$262 \Delta_f(x) = \operatorname{argmax}_{y \in (\{x\} \cup \{x' \mid (f(x') - f(x)) > c(x, x')\})} (f(y)) \quad (1)$$

263 We will denote it by Δ when f is clear from the context. $\Delta(x)$ might not be well defined if
264 there are multiple values of y that attains the maximum. In those cases, $\Delta(x)$ is chosen to
265 be the smallest y amongst them. In words, you jump to another feature only if the cost of
266 jumping is less than the advantage in being classified as 1.

267 We define the Jury's utility for publishing f ($U(f)$) as the classification accuracy with
268 respect to $h(x)$. Thus, Jury's utility for publishing f is

$$269 U(f) = \sum_{x \in \mathcal{X}} \pi(x) [f(\Delta(x)) \cdot h(x) + (1 - f(\Delta(x)) \cdot (1 - h(x)))] \\ 270 = \sum_{x \in \mathcal{X}} \pi(x) [f(\Delta(x)) \cdot (2h(x) - 1) + 1 - h(x)] \\ 271$$

272 We define $C(f) = \sum_{x \in \mathcal{X}} \pi(x) [h(x) \cdot c(x, \Delta_f(x))]$ to be the cost of strategy for a published
273 classifier f .

274 We define the efficiency of the classifier f ($E(f)$)² as follows:

$$275 E(f) = U(f) - C(f) \\ 276 = \sum_{x \in \mathcal{X}} \pi(x) [f(\Delta(x)) \cdot h(x) + (1 - f(\Delta(x)) \cdot (1 - h(x))] - \sum_{x \in \mathcal{X}} \pi(x) [h(x) \cdot c(x, \Delta(x))] \\ 277 = \sum_{x \in \mathcal{X}} \pi(x) [f(\Delta(x)) \cdot h(x) + (1 - f(\Delta(x)) \cdot (1 - h(x)) - h(x) \cdot c(x, \Delta(x))] \\ 278$$

279 The focus of this paper is to demonstrate what role randomness and noise can play in
280 strategic classification and not to give algorithms for learning the optimal or most efficient
281 strategic classifier. We can present the ideas even by making the following assumptions on
282 the cost function $c : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$:

- 283 1. $c(x, x') \geq 0, \forall x, x' \in \mathcal{X}$.
- 284 2. $c(x', x) = 0, \forall x, x' \mid h(x') \geq h(x)$, that is, jumping to a lesser qualified feature is free.
- 285 3. $c(x, x'') \leq c(x, x') + c(x', x''), \forall x, x', x'' \in \mathcal{X}$, that is, the costs are sub-additive.
- 286 4. $c(x, x') \leq c(x, x''), \forall x, x', x'' \mid h(x'') \geq h(x')$, that is, jumping to a lesser qualified feature
287 is easier.
- 288 5. $c(x', x'') \leq c(x, x''), \forall x, x', x'' \mid h(x') \geq h(x)$, that is, jumping from a lesser qualified
289 feature is harder.

290 The last two points are implied by the first three, we wrote them as separate points for
291 completeness. We call the cost function *simple* if it satisfies all the above assumptions.

292 By the virtue of the definition of simple cost functions, without loss of generality, we
293 assume that h is monotonically increasing with the feature x , that is, $\forall x, x' \in \mathcal{X}, x' \geq x \implies h(x') \geq h(x)$.

294 Next, we mention a special kind of cost function that satisfies the assumptions: $c(x, x') = \max(a(x') - a(x), 0)$ where the function $a : \mathcal{X} \rightarrow \mathbb{R}$ is monotonically increasing in x , that is,
295 $x' \geq x \implies a(x') \geq a(x)$.

¹ Such a best response model has been studied in the literature, for example, [17].

² We defined efficiency as $U(f) - C(f)$ for the simplicity of the presentation. Defining efficiency as $U(f) - \beta \cdot C(f)$ (for some $\beta > 0$) doesn't effect the theorems except for Theorem 3, which is no longer true for $\beta < 1$.

298 Given a cost function c , let

299 $\text{Lip}_1(c) = \{f \mid f : \mathcal{X} \rightarrow [0, 1], f(x') - f(x) \leq c(x, x') \forall x, x' \in \mathcal{X}\}$

300 Given the cost function c , we say f satisfies the Lipschitz constraint if $f \in \text{Lip}_1(c)$. Note
 301 that any classifier $f \in \text{Lip}_1(c)$ is monotonically increasing with x , that is, $x' \geq x \implies$
 302 $f(x') \geq f(x)$. This is because $\forall x' \geq x, f(x) - f(x') \leq c(x', x) = 0$. And $\forall x \in \mathcal{X}, \Delta_f(x) = x$,
 303 that is, no one changes their feature if f is the published classifier.

304 In Section 5, we generalize this model to the setting where the features are a noisy
 305 representation of an individual's private signal. An individual can make efforts to change
 306 their private signal but can't control the noise. The Jury only see the features and classifies
 307 an individual based on that. In Section 5, the fairness notion, we will concern ourselves with,
 308 is the classification accuracy of the published classifier across subpopulations.

309 **3 Committed Randomness Helps both Utility and Efficiency**

310 In this section, we compare the optimal utility and efficiency achieved by a deterministic
 311 binary classifier to a probabilistic classifier. Consider the following two scenarios:

312 *Scenario 1:* The Jury commits to using a binary classifier $f : \mathcal{X} \rightarrow \{0, 1\}$. The best
 313 response function $\Delta_f : \mathcal{X} \rightarrow \mathcal{X}$, Jury's utility from publishing f ($U(f)$) and efficiency of the
 314 classifier f ($E(f)$) are defined as in Section 2.

315 *Scenario 2:* The Jury publishes a probabilistic classifier $f : \mathcal{X} \rightarrow [0, 1]$ and commits to
 316 it. The best response function $\Delta_f : \mathcal{X} \rightarrow \mathcal{X}$, Jury's utility from publishing f ($U(f)$) and
 317 efficiency of the classifier f ($E(f)$) are as defined in Section 2. Note that this is equivalent to
 318 when Jury publishes a list of deterministic classifiers and chooses a classifier uniformly at
 319 random from them. Contestants update their feature without knowing which classifier gets
 320 picked up at the end.

321 The following example illustrates how randomization helps in getting strictly better utility
 322 and efficiency:

Let $\mathcal{X} = \{1, 2\}$ and each feature contains half of the population. Let

$$h(x) = \begin{cases} 1 & \text{if } x = 2 \\ 0 & \text{otherwise} \end{cases}$$

Let the cost of changing the feature from 1 to 2 be 0.5. The the randomized classifier f defined as follows:

$$f(x) = \begin{cases} 1 & \text{if } x = 2 \\ 0.5 & \text{if } x = 1 \end{cases}$$

323 achieves an accuracy of 0.75. The contestants at $x = 2$ are happy as they are already being
 324 classified as 1 with probability 1. For the contestants at $x = 1$, $f(2) - f(1) = 0.5 = c(1, 2)$
 325 and hence, they don't have an incentive to manipulate their feature. As all the contestants
 326 retain their true features, the efficiency of f is also equal to 0.75. As the feature space is
 327 bounded, there are only three options for a deterministic classifier: keep the threshold at 1
 328 and classify everyone as 1; keep the threshold at 2 and you end up classifying everyone as 1,
 329 as the contestants at 1 change their feature to 2; classify everyone as 0. All these classifiers
 330 have 0.5 accuracy and at most 0.5 efficiency.

331 In the mathematical example given above, the randomized classifier was set up such that
 332 none of the contestants had any incentive to change their feature. In the next subsection,
 333 we show that the most efficient classifier always looks like "this" for "simple" cost functions.

334 That is, if the cost function c satisfies the assumptions made in Section 2, then for every
 335 true qualification function h , there exists a function $f_h \in \text{Lip}_1(c)$ that achieves the optimal
 336 efficiency.

337 **3.1 Most Efficient Classifier for Simple Cost Functions**

338 Recall, $E(f) = \sum_{x \in \mathcal{X}} \pi(x)[f(\Delta(x)) \cdot h(x) + (1 - f(\Delta(x))) \cdot (1 - h(x)) - h(x) \cdot c(x, \Delta(x))]$. Let
 339 $E^* = \max_{f: \mathcal{X} \rightarrow [0,1]} \sum_{x \in \mathcal{X}} \pi(x)[f(\Delta(x)) \cdot h(x) + (1 - f(\Delta(x))) \cdot (1 - h(x)) - h(x) \cdot c(x, \Delta(x))]$.

340 **Theorem 3.** *For every monotone true qualification function $h: \mathcal{X} \rightarrow [0, 1]$, probability
 341 distribution $\pi: \mathcal{X} \rightarrow [0, 1]$ over the features, simple cost function c , there exists $g \in \text{Lip}_1(c)$
 342 such that $E(g) = E^*$.*

Proof. Let f be an efficiency maximizing classifier. We argue that $g: \mathcal{X} \rightarrow [0, 1]$ defined as

$$g(x) = \max_y \{f(y) - c(x, y)\}$$

343 is in $\text{Lip}_1(c)$ and satisfies $E(g) \geq E(f)$. Let $\delta_f(x) = \text{argmax}_y \{f(y) - c(x, y)\}$. When f is
 344 clear from the context, we will drop the subscript on δ . Using definition of δ , $g(x) \in [0, 1]$ as
 345 $\forall x, y \in \mathcal{X}, f(y) - c(x, y) \leq f(y) \leq 1$ ($c(x, y) \geq 0$) and $\max_y \{f(y) - c(x, y)\} \geq f(x) - c(x, x) \geq$
 346 0. For all $x, x' \in \mathcal{X}$,

$$\begin{aligned} 347 \quad g(x') - g(x) &= f(\delta(x')) - c(x', \delta(x')) - f(\delta(x)) + c(x, \delta(x)) \\ 348 \quad &= f(\delta(x')) - c(x, \delta(x')) - f(\delta(x)) + c(x, \delta(x)) + (c(x, \delta(x')) - c(x', \delta(x'))) \\ 349 \quad &\leq c(x, \delta(x')) - c(x', \delta(x')) \leq c(x, x') \quad (\text{sub-additivity}) \end{aligned}$$

351 The first inequality follows the definition of δ , that is, $\forall y \in \mathcal{X}, f(\delta(x)) - c(x, \delta(x)) \geq$
 352 $f(y) - c(x, y)$. Therefore, $f(\delta(x')) - c(x, \delta(x')) - f(\delta(x)) + c(x, \delta(x)) \leq 0$. The second
 353 inequality follows from the fact that the cost function c is simple and satisfies the sub-
 354 additivity condition. This proves that $g \in \text{Lip}_1(c)$. This implies, as observed previously,
 355 $\forall x \in \mathcal{X}, \Delta_g(x) = x$. Next, we show that $E(g) \geq E(f)$ and hence $E(g) = E^*$. Efficiency of
 356 the classifier g is

$$\begin{aligned} 357 \quad E(g) &= \sum_{x \in \mathcal{X}} \pi(x)[g(\Delta_g(x)) \cdot h(x) + (1 - g(\Delta_g(x))) \cdot (1 - h(x)) - h(x) \cdot c(x, \Delta_g(x))] \\ 358 \quad &= \sum_{x \in \mathcal{X}} \pi(x)[2 \cdot g(x) \cdot h(x) - g(x) - h(x) + 1] \end{aligned}$$

360 Efficiency of the classifier f is

$$\begin{aligned} 361 \quad E(f) &= \sum_{x \in \mathcal{X}} \pi(x)[f(\Delta_f(x)) \cdot h(x) + (1 - f(\Delta_f(x))) \cdot (1 - h(x)) - h(x) \cdot c(x, \Delta_f(x))] \\ 362 \quad &= \sum_{x \in \mathcal{X}} \pi(x)[2f(\Delta(x)) \cdot h(x) - f(\Delta(x)) - h(x) + 1 - h(x) \cdot c(x, \Delta(x))] \\ 364 \quad E(g) - E(f) &= \sum_{x \in \mathcal{X}} \pi(x)[(g(x) - f(\Delta(x))) \cdot (2h(x) - 1) + h(x) \cdot c(x, \Delta(x))] \\ 365 \quad \triangleright \text{Claim 4. } \forall x, [(g(x) - f(\Delta(x))) \cdot (2h(x) - 1) + h(x) \cdot c(x, \Delta(x))] &\geq 0. \end{aligned}$$

366 Please refer to Appendix A for the proof of the claim. It's straightforward to see that
 367 $E(g) - E(f) \geq 0$ using the above claim. Therefore, we showed a classifier $g \in \text{Lip}_1(c)$ such
 368 that $E(g) = E^*$. ◀

369 In words, *when we are concerned with the efficiency of the published classifier, the optimal is
 370 achieved by a probabilistic classifier that has zero cost of strategy and gives individuals no
 371 incentive to change their feature.*

4 Are Randomized Classifiers in Equilibrium from Jury's Perspective?

As discussed in the Section 1, there are many obstacles to implementing a randomized classifier in the strategic setting. In this section, we illustrate the instability caused by the use of randomized classifiers (which becomes increasingly important while considering multiple classifiers). In Section 3, we saw that a randomized classifier can achieve better accuracy and efficiency than any binary classifier. While maximizing efficiency, we further showed that the optimally efficient classifier is such that every contestant reveals their true feature. Once the Jury knows the contestants' true features, she can be greedy and classify the individuals using a threshold function with $\tau = \min\{x \mid h(x) \geq \frac{1}{2}\}$ as the threshold to achieve the best accuracy. Therefore, unless the Jury commits to using randomness, she has an incentive of not sticking to the promised randomized classifier. The question is: what's the best accuracy/efficiency achieved by a classifier that is in equilibrium even from Jury's perspective? We formalize this equilibrium concept as follows (the true qualification function h and the cost function c are fixed):

1. Jury publishes a randomized classifier $f : \mathcal{X} \rightarrow [0, 1]$.
2. Contestants, knowing f , changes their feature from x to $\Delta_f(x)$.
3. f is in equilibrium from Jury's perspective if given that the contestants changed their features according to the best response function Δ_f , f achieves the best classification accuracy, that is, for all classifiers $g \in \mathcal{X} \rightarrow [0, 1]$,

$$\sum_{x \in \mathcal{X}} \pi(x)[f(\Delta_f(x)) \cdot h(x) + (1 - f(\Delta_f(x))) \cdot (1 - h(x))] \quad (2)$$

$$- \sum_{x \in \mathcal{X}} \pi(x)[g(\Delta_f(x)) \cdot h(x) + (1 - g(\Delta_f(x))) \cdot (1 - h(x))] \geq 0$$

Using next theorem, we show that for any randomized classifier that is in equilibrium from Jury's perspective, there exists a binary classifier that achieves at least the same accuracy.

► **Theorem 5.** *Given a monotone true qualification function h , probability distribution π over the features, and a simple cost function c , let $f^* : \mathcal{X} \rightarrow \{0, 1\}$ be the classifier that optimizes Jury's utility over the deterministic classifiers under Stackelberg equilibrium. Let $f : \mathcal{X} \rightarrow [0, 1]$ be a randomized classifier such that $U(f) > U(f^*)$, then f is not in an equilibrium from Jury's perspective (the notion defined above).*

Please refer to Appendix B for the proof.

Disclaimer: f' as defined above might also not be in equilibrium from Jury's perspective. The above theorem illustrates the following point: *Jury doesn't benefit from randomized classifiers without creating instability in the system.*

Can we somehow exploit this power of randomness while overcoming the obstacles to randomized classification? The answer is yes – make the features noisy.

5 Noisy Features Give the System Free Randomness

We formalize the setting with noisy features as follows: every individual has a private signal $y \in \mathcal{X}$. The true qualification function $h : \mathcal{X} \rightarrow [0, 1]$ depends on y , that is, $h(y)$ is the probability of an individual being qualified (1) given that its private signal is y . Given a private signal y , a feature is drawn randomly from the distribution $p_y : \mathcal{X} \rightarrow [0, 1]$, that is, $p_y(x)$ is the probability that an individual's feature is x when their private signal is y . If $\mathcal{X} = \mathbb{R}$, the right intuition for p_y is it being $\mathcal{N}(y, \sigma)$ where $\mathcal{N}(y, \sigma)$ is the gaussian

414 distribution with mean y and standard deviation σ . Let $\pi : \mathcal{X} \rightarrow [0, 1]$ be the probability
 415 distribution over the private signals y realized by the individuals.

416 Let $c(y, y')$ be the cost incurred by the contestant to change their private signal from y
 417 to y' . The contestants can put effort to change their private signals but the feature would
 418 still be drawn randomly using the updated private signal.

419 The classification is again modeled as a sequential game where a Jury publishes a
 420 deterministic classifier $f : \mathcal{X} \rightarrow \{0, 1\}$. We restricts ourselves to deterministic classifiers due
 421 to the observations made in Section 4. Contestants change their private signals as long as
 422 they are ready to incur the cost of change. Given a private signal y , let $q_f(y)$ denote the
 423 probability of a contestant, with private signal y , being classified as 1 when f is the classifier.
 424 Therefore, $q_f(y) = \sum_{x \in \mathcal{X}} p_y(x) \cdot f(x)$.

425 Given f , the best response of a contestant with private signal y is given as,

$$426 \quad \Delta_f(y) = \operatorname{argmax}_{z \in \{y\} \cup \{y' | q_f(y') - q_f(y) > c(y, y')\}} (q_f(z)) \quad (3)$$

427 We will denote it by Δ when f is clear from the context. $\Delta(y)$ might not be well defined
 428 if there are multiple values of z that attains the maximum. In those cases, $\Delta(y)$ is chosen
 429 to be the smallest z amongst them. In words, you jump to another private signal only if
 430 the cost of jumping is less than the advantage in being classified as 1. Even though f is
 431 deterministic, due to noisy features, the effective classifier given the private signal y (q_f) is
 432 probabilistic. Therefore, we will see below that the noise allows us similar advantages as that
 433 of a probabilistic classifier.

The accuracy and efficiency of the classifier f are defined as follows:

$$U(f) = \sum_{y \in \mathcal{X}} \pi(y) [q_f(\Delta(y)) \cdot h(y) + (1 - q_f(\Delta(y))) \cdot (1 - h(y))]$$

$$E(f) = \sum_{y \in \mathcal{X}} \pi(y) [q_f(\Delta(y)) \cdot h(y) + (1 - q_f(\Delta(y))) \cdot (1 - h(y))] - \sum_{y \in \mathcal{X}} \pi(y) [h(y) \cdot c(y, \Delta(y))]$$

434 We assume that h is monotonically increasing with y and the cost function c is simple. Next,
 435 we will demonstrate how noisy features can lead fairer outcomes and even increase Jury's
 436 accuracy.

437 5.1 Noisy Features achieve Fairer Equilibria

438 Consider two subpopulations A and B . For simplicity, these subpopulations are a partition
 439 of the individuals in the universe. Let s_A denote the probability an individual from the
 440 universe is in subpopulation A . Similarly, s_B ($s_A = 1 - s_B$). Let $h_A : A \rightarrow [0, 1]$ be the true
 441 qualification function for the subpopulation A . Similarly, h_B . Let $c_A : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ be the
 442 cost function for the subpopulation A , that is, $c_A(y, y')$ is the cost of changing the private
 443 signal from y to y' for an individual in A . Similarly, c_B is defined. Let $\pi_A : A \rightarrow [0, 1]$ and
 444 π_B be the probability distribution over the private signals realized by the subpopulations A
 445 and B respectively.

446 Given a published deterministic classifier $f : \mathcal{X} \rightarrow \{0, 1\}$, the best response of the
 447 contestant in subpopulation A with private signal y ($\Delta_f^A(y)$) is defined using c_A as the
 448 cost function. Similarly, for subpopulation B , let $\Delta_f^B(y)$ denote the best response of the
 449 contestant in subpopulation B with private signal y and when the published classifier is
 450 f . We use $U_A(f)$ and $U_B(f)$ to denote the accuracy of the classifier f on the respective
 451 subpopulations.

9:12 The Role of Randomness and Noise in Strategic Classification

452 We consider the setting where $h_A = h_B = h$ and $\pi_A = \pi_B = \Pi$, but the cost functions
 453 c_A and c_B are different. In this section, we use the symbol Π to denote the probability
 454 distribution over the private signals to avoid confusion with the Archimedes' constant π .

455 In our first example, we show that even though the subpopulations are identical with
 456 respect to their qualifications, different costs can lead to unfair classification when classification
 457 is based on private signals. Through our second example, we show that the use of noisy
 458 features, for strategic classification, can lead to increase in the overall accuracy of classification
 459 as well as give fair classification. We evaluate the fairness of a classifier f quantitatively using
 460 the difference between the accuracies, that is, $|U_A(f) - U_B(f)|$.

461 Let's start with the example. $\mathcal{X} = \mathbb{R}$. Let the true qualification function for both the

462 subpopulations be as follows: $h(y) = \begin{cases} 1 & \text{if } y > d \\ \frac{y}{2d} + \frac{1}{2} & \text{if } y \in [-d, d] \\ 0 & \text{if } y < -d \end{cases}$, where d is a fixed large
 463 enough positive real number. Let the probability density function on the private signals
 464 realized by the subpopulations be as follows: $\Pi(y) = \frac{e^{-\frac{y^2}{2t^2}}}{\sqrt{2\pi}t}$, that is, the gaussian distribution
 465 with mean 0 and standard deviation t . Again, t is fixed positive real number. We assume
 466 $d \gg t$.

467 Let σ_A and σ_B be positive real numbers. The cost function for a subpopulation $S \in \{A, B\}$
 468 is defined as follows (with $(y' - y)^+ = \max\{y' - y, 0\}$):

$$469 \quad c_S(y, y') = \frac{(y' - y)^+}{\sqrt{2\pi}\sigma_S} \quad (4)$$

470 We start with the setting where the features are the private signals and not a noisy
 471 representation of them.

472 *Remark:* If the Jury is allowed to publish different classifiers for the two subpopulations,
 473 then she can achieve "the best possible accuracy" on both the subpopulations. It's easy to
 474 see that the classifier $f_S : \mathcal{X} \rightarrow \{0, 1\}$, defined as follows, achieves as much accuracy as a
 475 classifier under no strategic manipulation of the features can achieve on the subpopulation

$$476 \quad S \in \{A, B\}: f_S(y) = \begin{cases} 1 & \text{if } y \geq \sqrt{2\pi}\sigma_S \\ 0 & \text{otherwise} \end{cases}$$

477 All the contestants in a subpopulation S , with $0 < y < \sqrt{2\pi}\sigma_S$ report their private signals
 478 to be $\sqrt{2\pi}\sigma_S$ as cost of this change is < 1 whereas the advantage gained in the probability of
 479 being classified as 1 is 1. For all the contestants with private signal $y \leq 0$, the cost of change
 480 is too high (≥ 1) and thus, they report their true private signals. Therefore, the classifier f_S
 481 ends up classifying everyone with private signal $y > 0$ as 1 which is the accuracy maximizing
 482 classification under the "no strategic manipulation" setting.

483 **How strategic classification leads to unfairness:** When $\sigma_A \neq \sigma_B$, the optimal
 484 classifiers for the subpopulations A and B are different and hence, when we choose a single
 485 classifier for both the subpopulations, we are bound to loose on the accuracy of at least one of
 486 the subpopulations. Through an example (Theorem 6), we suggest that: *while maximizing the*
 487 *overall accuracy over the universe, the minority group might be at a disadvantage irrespective*
 488 *of whether their costs to change the private signals are higher or lower than the majority*
 489 *subpopulation.* Without loss of generality, we assume that A is the minority subpopulation,
 490 that is, $s_A \leq s_B$. In many real life scenarios, the Jury would publish a single classifier for
 491 both the subpopulations either because A is a protected group and the Jury is not allowed to
 492 discriminate based on the subgroup membership or because the Jury has not yet identified
 493 these subpopulations and the differences in their cost functions.

494 ► **Theorem 6.** Let A and B be two subpopulations such that the true qualification functions,
 495 h_A, h_B , the probability density functions, π_A, π_B and the cost functions c_A, c_B are as
 496 instantiated above.

497 Assuming $|\sigma_A - \sigma_B| \leq \frac{t}{\sqrt{2\pi}}$, let f^* be the deterministic classifier that maximizes Jury's
 498 utility $U(f)$, if $s_A < s_B$ and $\sigma_A \neq \sigma_B$ (the cost functions are different), then $U_A(f^*) < U_B(f^*)$, that is, the minority is at a disadvantage, even though their qualifications were
 499 identical ($h_A = h_B, \pi_A = \pi_B$).

500 Please refer to Appendix C for the proof.

501 Next we show that, when the features are appropriately noisy, the optimal classifier from
 502 Jury's perspective is fair to the subpopulations. The intuition is as follows: if the noise is
 503 large enough such that none of contestants in either of the subpopulations want to manipulate
 504 their private signals, then the cost differences become irrelevant and hence, the optimal
 505 classifier achieves equal accuracy on both the subpopulations. You would think that this
 506 addition of noise would compromise Jury's utility. Subsequently, we show that adding noise
 507 might also improve the overall accuracy of the Jury's optimal classifier, therefore, addition of
 508 noise can make everyone happier. The latter is a continuation to the results at the start of
 509 Section 5 about the usefulness of noise to the Jury under strategic classification.

510 **Noisy features lead to fairer outcomes:** Now, we analyze the setting with noisy
 511 features and prove the following theorem. The true qualification function h , cost functions
 512 (c_A and c_B) and the probability density function Π are as defined for the first example. Let
 513 $\sigma = \max\{\sigma_A, \sigma_B\}$. Given a private signal y , the features x are distributed according to the
 514 gaussian with mean y and standard deviation σ . The probability density function for the
 515 feature x given the private signal y is $p_y(x) = \frac{e^{-\frac{(x-y)^2}{2\sigma^2}}}{\sqrt{2\pi\sigma}}$.

516 ► **Theorem 7.** Let A and B be two subpopulations such that the true qualification functions,
 517 h_A, h_B , the probability density functions, π_A, π_B and the cost functions c_A, c_B are as
 518 instantiated above. When the features are drawn with a gaussian noise of mean 0 and standard
 519 deviation σ , such that, $\sigma \geq \sigma_A, \sigma_B$, if f^* is the deterministic classifier that maximizes Jury's
 520 utility $U(f)$, then f^* is fair, that is, $U_A(f^*) = U_B(f^*)$.

521 Please refer to Appendix D for the proof.

522 Theorem 7 would hold for when we are concerned with multiple subpopulations as long as
 523 $\sigma \geq \sigma_S$ for every relevant subpopulation S . In words, using noisy features we *can* ensure that
 524 the best response of a Jury, maximizing her own utility, is fair to all the subpopulations that
 525 are identical in terms of qualifications but different in terms of the costs to manipulate the
 526 private signals, as long as the costs of manipulation for a subpopulation are not too small.

527 **Noisy features can also improve Jury's utility:** Next, we show that further in some
 528 cases, the addition of noise to the features is not only beneficial for ensuring fairness but
 529 might also achieve better overall accuracy under strategic classification compared to when a
 530 noiseless signal is used.

531 Retaining the instantiations of $h_A, h_B, \pi_A, \pi_B, c_A, c_B$ and σ as above, consider the
 532 following two scenarios: 1. Jury bases her classifier on the private signal y . 2. The features
 533 are drawn with a gaussian noise of mean 0 and standard deviation σ and Jury bases her
 534 classifier on the features (x) .

535 Let f_0^* and f_σ^* be the optimal classifiers under strategic classification in the two scenarios
 536 respectively. Let $U(f_0^*)$ be the overall classification accuracy (Jury's utility) under Scenario 1
 537 and $U(f_\sigma^*)$ be the overall classification accuracy (Jury's utility) under Scenario 2. We assume

540 that the subpopulations are equally populated, that is, $s_A = s_B$ for simplicity of calculations
 541 in the next theorem.

542 ► **Theorem 8.** *There exists qualification functions, h_A, h_B , the probability density functions
 543 over the private signals, π_A, π_B , the cost functions c_A, c_B and $\sigma > 0$ such that, $U(f_\sigma^*) >$
 544 $U(f_0^*)$, that is, the Jury gets better classification accuracy when the features are drawn with a
 545 gaussian noise of mean 0 and standard deviation σ . Here, the subpopulations have identical
 546 qualifications ($h_A = h_B, \pi_A = \pi_B$) but different cost functions.*

547 Please refer to Appendix E for the proof. This theorem corroborates the idea that not only
 548 the subpopulations, but even the Jury might prefer noisy features. In the above example,
 549 for simplicity, we assumed $s_A = s_B$. Therefore, the optimal classifier was fair even in the
 550 noiseless setting. But a slight tweak in s_A so that $s_A < s_B$ wouldn't change Jury's utility, in
 551 Scenario 1, by much and thus, would give an example where the noiseless setting has both
 552 unfairness and lesser overall classification accuracy.

553 In this paper, we study the interaction of noise with strategic classification through some
 554 simple examples, and leave the task of generalizing these results for future research.

555 6 Discussion

556 The problem of classification (and the strategic classification problem it entails) is of tremen-
 557 dous importance both practically (affecting pretty much every industry) and theoretically
 558 (with implications ranging from algorithms to policy and law). Therefore, clarifying the
 559 role randomness plays in this specific family of games is an important goal. Just as in
 560 games, randomness may lead to better solution in strategic classification. Moreover, in many
 561 important settings (such as college admissions in some jurisdictions), the classifier is required
 562 to be deterministic by law — which is not a handicap for algorithmic classification, but
 563 is a handicap for strategic one. In addition, we proved that, in many natural cases, any
 564 randomized classifier (based on one-dimension) that achieves strictly better accuracy than
 565 the optimal deterministic one is not stable from the classifier's standpoint, thus illustrating
 566 the difficulty of implementing a randomized classifier in a more complicated scenario with
 567 multiple classifiers (such as college admissions). This motivates the use of noisy features as a
 568 commitment device, which can improve both accuracy and fairness, and is also practically
 569 possible (for example by restricting the types of information available to the classifier).

570 References

- 571 1 Emrah Akyol, Cedric Langbort, and Tamer Basar. Price of transparency in strategic machine
 572 learning. *arXiv preprint arXiv:1610.08210*, 2016.
- 573 2 Michael Brückner and Tobias Scheffer. Stackelberg games for adversarial prediction problems.
 574 In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery
 575 and data mining*, pages 547–555. ACM, 2011.
- 576 3 Yiling Chen, Chara Podimata, Ariel D Procaccia, and Nisarg Shah. Strategyproof linear
 577 regression in high dimensions. In *Proceedings of the 2018 ACM Conference on Economics and
 578 Computation*, pages 9–26. ACM, 2018.
- 579 4 Jinshuo Dong, Aaron Roth, Zachary Schutzman, Bo Waggoner, and Zhiwei Steven Wu.
 580 Strategic classification from revealed preferences. In *Proceedings of the 2018 ACM Conference
 581 on Economics and Computation*, pages 55–70. ACM, 2018.
- 582 5 Kfir Eliaz and Ran Spiegler. The model selection curse. *American Economic Review: Insights*.
- 583 6 Richard Engelbrecht-Wiggans. On the value of private information in an auction: ignorance
 584 may be bliss. *BEBR faculty working paper; no. 1242*, 1986.

585 7 Alex Frankel and Navin Kartik. Improving information from manipulable data. *arXiv preprint*
 586 *arXiv:1908.10330*, 2019.

587 8 Moritz Hardt, Nimrod Megiddo, Christos Papadimitriou, and Mary Wootters. Strategic
 588 classification. In *Proceedings of the 2016 ACM conference on innovations in theoretical*
 589 *computer science*, pages 111–122. ACM, 2016.

590 9 Lily Hu, Nicole Immorlica, and Jennifer Wortman Vaughan. The disparate effects of strategic
 591 manipulation. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*,
 592 pages 259–268. ACM, 2019.

593 10 Nicole Immorlica, Katrina Ligett, and Juba Ziani. Access to population-level signaling as
 594 a source of inequality. In *Proceedings of the Conference on Fairness, Accountability, and*
 595 *Transparency*, pages 249–258. ACM, 2019.

596 11 Sampath Kannan, Aaron Roth, and Juba Ziani. Downstream effects of affirmative action. In
 597 *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pages 240–248.
 598 ACM, 2019.

599 12 Andrew Kephart and Vincent Conitzer. Complexity of mechanism design with signaling costs.
 600 In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent*
 601 *Systems*, pages 357–365, 2015.

602 13 Andrew Kephart and Vincent Conitzer. The revelation principle for mechanism design with
 603 reporting costs. In *Proceedings of the 2016 ACM Conference on Economics and Computation*,
 604 pages 85–102, 2016.

605 14 Jon Kleinberg and Manish Raghavan. How do classifiers induce agents to invest effort
 606 strategically? In *Proceedings of the 2019 ACM Conference on Economics and Computation*,
 607 pages 825–844. ACM, 2019.

608 15 John Miller, Smitha Milli, and Moritz Hardt. Strategic adaptation to classifiers: A causal
 609 perspective. *arXiv preprint arXiv:1910.10362*, 2019.

610 16 Smitha Milli, John Miller, Anca D Dragan, and Moritz Hardt. The social cost of strategic
 611 classification. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*,
 612 pages 230–239. ACM, 2019.

613 17 Christopher A Wilkens, Ruggiero Cavallo, Rad Niazadeh, and Samuel Taggart. Mechanism
 614 design for value maximizers. *arXiv preprint arXiv:1607.04362*, 2016.

615 **A Proof of Claim 4**

616 Recalling, $g(x) = f(\delta(x)) - c(x, \delta(x))$. Using definition of δ , we know that

617
$$g(x) = f(\delta(x)) - c(x, \delta(x)) \geq f(\Delta(x)) - c(x, \Delta(x)) \quad (5)$$

619 And, using definition of Δ , we can show that

620
$$f(\Delta(x)) \geq g(x) \quad (6)$$

622 This is because, either $f(\delta(x)) - c(x, \delta(x)) = f(x)$ and as $f(\Delta(x)) \geq f(x)$, we get the
 623 inequality. Or, $f(\delta(x)) - c(x, \delta(x)) > f(x)$, which implies that x has an incentive to change its
 624 feature to $\delta(x)$. Therefore, by the definition of Δ , $f(\Delta(x)) \geq f(\delta(x)) \geq f(\delta(x)) - c(x, \delta(x))$.
 625 The expression in the claim can be rewritten as

626
$$(g(x) - f(\Delta(x))) \cdot (2h(x) - 1) + h(x) \cdot c(x, \Delta(x))$$

 627
$$= (g(x) - f(\Delta(x))) \cdot (h(x) - 1) + h(x) \cdot (g(x) - f(\Delta(x)) + c(x, \Delta(x)))$$

629 As $g(x) - f(\Delta(x)) \leq 0$ from Equation 6 and $g(x) - f(\Delta(x)) + c(x, \Delta(x)) \geq 0$ from Equation
 630 5, the inequality follows from the fact that $0 \leq h(x) \leq 1$. This proves the claim.

631 **B Proof of Theorem 5**

 632 Equation 2 implies that for all classifiers $g \in \mathcal{X} \rightarrow [0, 1]$,

$$\begin{aligned} 633 \quad & \sum_{x \in \mathcal{X}} \pi(x) [(f(\Delta_f(x)) - g(\Delta_f(x))) \cdot (2h(x) - 1)] \geq 0 \\ 634 \quad \implies & \sum_{y \in \mathcal{X}} (f(y) - g(y)) \cdot \sum_{x: \Delta_f(x)=y} \pi(x) (2h(x) - 1) \geq 0 \\ 635 \end{aligned}$$

 636 Therefore, if f is in equilibrium from the Jury's perspective, for all $y \in \mathcal{X}$ such that
 637 $f(y) \in (0, 1)$, $\sum_{x: \Delta_f(x)=y} \pi(x) (2h(x) - 1) = 0$ otherwise Jury can choose $g(y) = 1$ (or 0)
 638 depending on whether $\sum_{x: \Delta_f(x)=y} \pi(x) (2h(x) - 1) > 0$ (or < 0) to increase her accuracy.
 639 Therefore, accuracy of the classifier f is given by

$$\begin{aligned} 640 \quad U(f) &= \sum_{x \in \mathcal{X}} \pi(x) [f(\Delta_f(x)) \cdot (2h(x) - 1) + (1 - h(x))] \\ 641 \quad &= \sum_{y \in \mathcal{X}} f(y) \cdot \sum_{x: \Delta_f(x)=y} \pi(x) (2h(x) - 1) + \sum_{x \in \mathcal{X}} \pi(x) (1 - h(x)) \\ 642 \quad &= \sum_{y: f(y)=1} \sum_{x: \Delta_f(x)=y} \pi(x) (2h(x) - 1) + \sum_{x \in \mathcal{X}} \pi(x) (1 - h(x)) \\ 643 \end{aligned}$$

 644 Consider a binary classifier $f' : \mathcal{X} \rightarrow \{0, 1\}$ defined as follows: $f(x) \in [0, 1] \implies f'(x) = 0$
 645 and $f(x) = 1 \implies f'(x) = 1$. We can show that $U(f') \geq U(f)$. The contestants who change
 646 their features when f' is the published classifier is a subset of $\{x \in \mathcal{X} \mid f(\Delta_f(x)) \in (0, 1)\}$
 647 and as $\sum_{x: f(\Delta_f(x)) \in (0, 1)} \pi(x) (2h(x) - 1) = 0$, the accuracy of f' can only increase. This is
 648 because: $\forall x \in \mathcal{X}$ if $f(\Delta_f(x)) = 0$, then $f'(\Delta_{f'}(x)) = 0$ as otherwise if x changed its feature
 649 under f' , it had an incentive to change under f too.

 650 If $x' > x$, $f(\Delta_f(x'))$, $f(\Delta_f(x)) \in (0, 1)$ and x changes its feature under f' , then x' has the
 651 incentive to change too as $c(x', x) = 0$, and hence, the subset of $\{x \in \mathcal{X} \mid f(\Delta_f(x)) \in (0, 1)\}$
 652 that change their features under f' can only do a positive addition to the utility (h is
 653 monotonically increasing with x and $\sum_{x: f(\Delta_f(x)) \in (0, 1)} \pi(x) (2h(x) - 1) = 0$). And, the
 654 contestants (x) who changed their features under f such that $f(\Delta_f(x)) = 1$ would also
 655 change their features under f' such that $f'(\Delta_{f'}(x)) = 1$ (as $f'(x) \leq f(x)$) and are already
 656 included in the calculation of $U(f)$.

 657 **C Proof of Theorem 6**

 Jury publishes a deterministic classifier and as there's no noise involved, without loss of generality, we can assume that f is a threshold classifier on the space \mathcal{X} (as c_A and c_B are simple cost functions). This assumption is justified in Section 3. Given the classifier $f : \mathcal{X} \rightarrow \{0, 1\}$ with threshold τ , the best response of a contestant in the subpopulation $S \in \{A, B\}$ is given as follows:

$$\Delta_f^S(y) = \begin{cases} y & \text{if } y \geq \tau \\ \tau & \text{if } \tau - \sqrt{2\pi}\sigma_S < y < \tau \\ y & \text{if } y \leq \tau - \sqrt{2\pi}\sigma_S \end{cases}$$

 658 The accuracy of the classifier f for the subpopulation S is given as follows:

$$\begin{aligned} 659 \quad U_S(f) &= \int_{-\infty}^{\infty} \Pi(y) [f(\Delta_f^S(y)) \cdot (2h(y) - 1) + (1 - h(y))] dy \\ 660 \end{aligned}$$

661 Let $c = \int_{-\infty}^{\infty} \Pi(y)[(1 - h(y))]dy$ which is independent of the subpopulation and the classifier.
 662 Therefore, $U_S(f) = \left(\int_{-\infty}^{\infty} \Pi(y)[f(\Delta^S(y)) \cdot (2h(y) - 1)]dy \right) + c$.

For the convenience of calculations, we will replace $h(y)$ with the following function,

$$h'(y) = \frac{y}{2d} + \frac{1}{2}$$

663 As d is large and Π is a gaussian centered at 0, this change barely affects the utility
 664 values. To be precise, the difference in the utility calculations for any classifier f while using
 665 h' instead of h is bounded by

$$\begin{aligned} 666 \quad \left| \int_{-\infty}^{\infty} \Pi(y)[f(\Delta^S(y)) \cdot 2(h(y) - h'(y))]dy \right| &\leq 2 \int_{-\infty}^{\infty} \Pi(y)[f(\Delta^S(y))|h(y) - h'(y)|]dy \\ 667 \quad &\leq 2 \int_{-\infty}^{\infty} \Pi(y) \cdot |h(y) - h'(y)|dy \\ 668 \quad &= 4 \int_d^{\infty} \Pi(y) \cdot \left(\frac{y}{2d} - \frac{1}{2} \right) dy \\ 669 \quad &\leq 2 \int_d^{\infty} \frac{e^{-\frac{y^2}{2t^2}}}{\sqrt{2\pi t}} \cdot \frac{y}{d} dy = 2 \frac{te^{-\frac{d^2}{2t^2}}}{\sqrt{2\pi d}} \end{aligned}$$

671 As we take d ($d \gg t$) to be large enough, we would be able to ignore this difference.
 672 From now onwards, we use h' as the "true qualification function".

673 Therefore, the accuracy of the classifier f over the subpopulation $S \in \{A, B\}$ can be
 674 approximated by

$$\begin{aligned} 675 \quad U_S(f) &= \left(\int_{-\infty}^{\infty} \Pi(y)[f(\Delta^S(y)) \cdot (2h'(y) - 1)]dy \right) + c = \left(\int_{-\infty}^{\infty} \Pi(y) \cdot f(\Delta^S(y)) \cdot \frac{y}{d} dy \right) + c \\ 676 \quad &= \left(\int_{\tau - \sqrt{2\pi}\sigma_S}^{\infty} \frac{e^{-\frac{y^2}{2t^2}}}{\sqrt{2\pi t}} \cdot \frac{y}{d} dy \right) + c = \frac{t}{\sqrt{2\pi d}} e^{-(\tau - \sqrt{2\pi}\sigma_S)^2/2t^2} + c \end{aligned}$$

678 The second last equality follows from the definition of Δ_f^S and the fact that f classifies
 679 everyone, with the updated private signal greater than or equal to τ , as 1 and 0 otherwise.

680

681 The overall accuracy of the classifier f is given by

$$\begin{aligned} 682 \quad U(f) &= s_A \cdot U_A(f) + s_B \cdot U_B(f) \\ 683 \quad &= s_A \cdot \frac{t}{\sqrt{2\pi d}} e^{-(\tau - \sqrt{2\pi}\sigma_A)^2/2t^2} + s_B \cdot \frac{t}{\sqrt{2\pi d}} e^{-(\tau - \sqrt{2\pi}\sigma_B)^2/2t^2} + c \end{aligned} \tag{7}$$

685 It's clear from the expression that the accuracy for the subpopulation A is maximized at
 686 $\tau_A = \sqrt{2\pi}\sigma_A$ and that of B is maximized at $\tau_B = \sqrt{2\pi}\sigma_B$. Consider the case when $s_A < s_B$.
 687 As $\tau_A \neq \tau_B$, and $U_B(f)$ has a larger weight in the expression, intuitively, while optimizing the
 688 overall accuracy, τ would try to achieve better accuracy for the subpopulation B , irrespective
 689 of whether $\sigma_A > \sigma_B$ or $\sigma_A < \sigma_B$, leading to unfairness across the subpopulations (A being
 690 at a disadvantage).

691 It's complicated to calculate the optimal τ , below we give a proof of the fact that the
 692 optimal τ would be such that $U_A(f) < U_B(f)$. To find the optimal value of τ , we differentiate
 693 $U(f)$ with respect τ as follows:

$$\begin{aligned}
 \frac{dU(f)}{d\tau} &= s_A \cdot \frac{dU_A(f)}{d\tau} + s_B \cdot \frac{dU_B(f)}{d\tau} \\
 &= -\frac{1}{\sqrt{2\pi}td} \left(s_A \cdot (\tau - \sqrt{2\pi}\sigma_A) \cdot e^{-(\tau - \sqrt{2\pi}\sigma_A)^2/2t^2} + s_B \cdot (\tau - \sqrt{2\pi}\sigma_B) \cdot e^{-(\tau - \sqrt{2\pi}\sigma_B)^2/2t^2} \right)
 \end{aligned}$$

697 Therefore, $\frac{dU(f)}{d\tau} = 0$

$$\begin{aligned}
 698 \quad & \implies s_A \cdot (\tau - \sqrt{2\pi}\sigma_A) \cdot e^{-(\tau - \sqrt{2\pi}\sigma_A)^2/2t^2} + s_B \cdot (\tau - \sqrt{2\pi}\sigma_B) \cdot e^{-(\tau - \sqrt{2\pi}\sigma_B)^2/2t^2} = 0 \\
 699 \quad & \implies \left| \frac{(\tau - \sqrt{2\pi}\sigma_A) \cdot e^{-(\tau - \sqrt{2\pi}\sigma_A)^2/2t^2}}{(\tau - \sqrt{2\pi}\sigma_B) \cdot e^{-(\tau - \sqrt{2\pi}\sigma_B)^2/2t^2}} \right| > 1 \quad (s_B > s_A) \\
 700
 \end{aligned}$$

701 As $ze^{-\frac{z^2}{2t^2}}$ is maximized at $z = t$, as long as $|\sigma_A - \sigma_B| \leq \frac{t}{\sqrt{2\pi}}$ (implying $|\tau - \sqrt{2\pi}\sigma_S| \leq t$ 702 for $S \in \{A, B\}$), the overall accuracy is maximized at a threshold τ such that $|\tau - \sqrt{2\pi}\sigma_A| > 703 |\tau - \sqrt{2\pi}\sigma_B|$ and hence, $U_A(f^*) < U_B(f^*)$, where f^* is the optimal classifier from Jury's 704 perspective. The assumption, $|\sigma_A - \sigma_B| \leq \frac{t}{\sqrt{2\pi}}$, can be interpreted as the subpopulations 705 being different but not extremely different, which is reasonable assumption in many real life 706 scenarios.

707 D Proof of Theorem 7

708 Again, we will replace the function h with h' (as in proof of Theorem 6) while loosing 709 an insignificant amount in all the calculations ($d >> t, \sigma$). Let $\Pi' : \mathcal{X} \rightarrow [0, 1]$ be the 710 probability density function over the features realized by each of the subpopulations. Let 711 $H(x)$ ($H : \mathcal{X} \rightarrow [0, 1]$) represent the probability of an individual being qualified (1) given that 712 the Jury sees feature x . These functions are same for both the subpopulations. As the Jury 713 only sees the feature and not the private signal, her accuracy is information-theoretically 714 limited by these functions as we will describe below. Firstly, $\Pi' : \mathcal{X} \rightarrow [0, 1]$ is given as 715 follows:

$$\begin{aligned}
 716 \quad \Pi'(x) &= \int_{-\infty}^{\infty} \Pi(y) \cdot p_y(x) dy = \int_{-\infty}^{\infty} \frac{e^{-\frac{y^2}{2t^2}}}{\sqrt{2\pi}t} \cdot \frac{e^{-\frac{(x-y)^2}{2\sigma^2}}}{\sqrt{2\pi}\sigma} dy \\
 717 \quad &= \int_{-\infty}^{\infty} \frac{e^{-\frac{x^2}{2(\sigma^2+t^2)}}}{\sqrt{2\pi}t} \cdot \frac{e^{-(y-\frac{x^2}{\sigma^2+t^2})^2/(2\frac{\sigma^2t^2}{\sigma^2+t^2})}}{\sqrt{2\pi}\sigma} dy \\
 718 \quad &= \frac{e^{-\frac{x^2}{2(\sigma^2+t^2)}}}{\sqrt{2\pi}t \cdot \sqrt{2\pi}\sigma} \int_{-\infty}^{\infty} e^{-(y-\frac{x^2}{\sigma^2+t^2})^2/(2\frac{\sigma^2t^2}{\sigma^2+t^2})} dy \\
 719 \quad &= \frac{e^{-\frac{x^2}{2(\sigma^2+t^2)}}}{\sqrt{2\pi}t \cdot \sqrt{2\pi}\sigma} \sqrt{2\pi \frac{\sigma^2t^2}{\sigma^2+t^2}} = \frac{e^{-\frac{x^2}{2(\sigma^2+t^2)}}}{\sqrt{2\pi(\sigma^2+t^2)}}
 \end{aligned}$$

721 Therefore, the probability density function over the features realized by the subpopulations, 722 with $\mathcal{N}(0, \sigma)$ gaussian noise, is itself a gaussian with mean 0 and $\sqrt{(\sigma^2 + t^2)}$ standard 723 deviation.

724 The qualification function given the features, H , is given as follows:

$$\begin{aligned}
 725 \quad H(x) &= \frac{1}{\Pi'(x)} \int_{-\infty}^{\infty} \Pi(y) \cdot p_y(x) \cdot h(y) dy \\
 726
 \end{aligned}$$

⁷²⁷ We replace h with h' , thus replacing H with H' as defined below:

$$\begin{aligned}
 H'(x) &= \frac{1}{\Pi'(x)} \int_{-\infty}^{\infty} \Pi(y) \cdot p_y(x) \cdot h'(y) dy = \frac{1}{\Pi'(x)} \int_{-\infty}^{\infty} \frac{e^{-\frac{y^2}{2t^2}}}{\sqrt{2\pi t}} \cdot \frac{e^{-\frac{(x-y)^2}{2\sigma^2}}}{\sqrt{2\pi\sigma}} \cdot \left(\frac{y}{2d} + \frac{1}{2}\right) dy \\
 &= \frac{1}{2} + \frac{1}{2 \cdot \Pi'(x)} \int_{-\infty}^{\infty} \frac{e^{-\frac{x^2}{2(\sigma^2+t^2)}}}{\sqrt{2\pi t}} \cdot \frac{e^{-(y-\frac{xt^2}{\sigma^2+t^2})^2/(2\frac{\sigma^2 t^2}{\sigma^2+t^2})}}{\sqrt{2\pi\sigma}} \cdot \frac{y}{2d} dy \\
 &= \frac{1}{2} + \frac{1}{2d \cdot \Pi'(x)} \frac{e^{-\frac{x^2}{2(\sigma^2+t^2)}}}{\sqrt{2\pi t} \cdot \sqrt{2\pi\sigma}} \int_{-\infty}^{\infty} e^{-(y-\frac{xt^2}{\sigma^2+t^2})^2/(2\frac{\sigma^2 t^2}{\sigma^2+t^2})} \cdot y dy \\
 &= \frac{1}{2} + \frac{1}{2d \cdot \Pi'(x)} \frac{e^{-\frac{x^2}{2(\sigma^2+t^2)}}}{\sqrt{2\pi t} \cdot \sqrt{2\pi\sigma}} \cdot \sqrt{2\pi \frac{\sigma^2 t^2}{\sigma^2+t^2}} \cdot \frac{xt^2}{\sigma^2+t^2} \\
 &= \frac{1}{2} + \frac{t^2}{\sigma^2+t^2} \frac{x}{2d}
 \end{aligned}$$

⁷³⁴ Therefore, when there's no strategic manipulation, Jury would classify any individual ⁷³⁵ with feature $x > 0$ as 1 and 0 otherwise. This is because, $H'(x) > \frac{1}{2}$ if and only if $x > 0$ and ⁷³⁶ the Jury would classify a feature as 1 if and only if, in expectation, the individuals with that ⁷³⁷ feature are more likely to be qualified. This is true irrespective of whether an individual ⁷³⁸ is from the subpopulation A or B because these subpopulations are identical in terms of ⁷³⁹ qualifications, that is, $h_A = h_B = h$ and $\pi_A = \pi_B = \Pi$.

⁷⁴⁰ We show that for the cost functions defined above, if Jury publishes $f = 1_{x>0}$, as the ⁷⁴¹ classifier, then none of the contestants in both the subpopulations A and B have an incentive ⁷⁴² to change their private signal (under $\mathcal{N}(0, \sigma)$ gaussian noise). Hence, the Jury gets the best ⁷⁴³ possible accuracy from these features and the classification is fair. For a subpopulation ⁷⁴⁴ $S \in \{A, B\}$, let $q_f^S(y)$ denote the probability of a contestant, with private signal y , being ⁷⁴⁵ classified as 1 when f is the classifier. Therefore,

$$q_f^S(y) = \int_{-\infty}^{\infty} f(x) \cdot p_y(x) dx = \int_0^{\infty} \frac{e^{-\frac{(x-y)^2}{2\sigma^2}}}{\sqrt{2\pi\sigma}} dx$$

⁷⁴⁸ For a subpopulation $S \in \{A, B\}$, let's calculate the advantage that a contestant, with ⁷⁴⁹ private signal y , gets by changing its signal to y' ($y' > y$, otherwise $q_f^S(y') \leq q_f^S(y)$):

$$\begin{aligned}
 q_f^S(y') - q_f^S(y) &= \int_0^{\infty} \frac{e^{-\frac{(x-y')^2}{2\sigma^2}}}{\sqrt{2\pi\sigma}} dx - \int_0^{\infty} \frac{e^{-\frac{(x-y)^2}{2\sigma^2}}}{\sqrt{2\pi\sigma}} dx = \int_{-y'}^{\infty} \frac{e^{-\frac{x^2}{2\sigma^2}}}{\sqrt{2\pi\sigma}} dx - \int_{-y}^{\infty} \frac{e^{-\frac{x^2}{2\sigma^2}}}{\sqrt{2\pi\sigma}} dx \\
 &= \int_{-y'}^{-y} \frac{e^{-\frac{x^2}{2\sigma^2}}}{\sqrt{2\pi\sigma}} dx \leq \int_{-y'}^{-y} \frac{1}{\sqrt{2\pi\sigma}} dx = \frac{y' - y}{\sqrt{2\pi\sigma}}
 \end{aligned}$$

⁷⁵³ As $\sigma = \max\{\sigma_A, \sigma_B\}$ and recalling the definitions of the cost functions c_A and c_B (Equation ⁷⁵⁴ 4), we get that

$$q_f^A(y') - q_f^A(y) \leq c_A(y, y') \quad \text{and} \quad q_f^B(y') - q_f^B(y) \leq c_B(y, y')$$

⁷⁵⁵ Therefore, none of the contestants in any of the subpopulations have an incentive to change

9:20 The Role of Randomness and Noise in Strategic Classification

758 their private signals. The accuracy of the classifier f on the subpopulation A is given as

$$\begin{aligned}
 759 \quad U_A(f) &= \left(\int_{-\infty}^{\infty} \Pi(y) [q_f^A(\Delta_f^A(y)) \cdot (2h(y) - 1)] dy \right) + c \\
 760 &= \left(\int_{-\infty}^{\infty} \Pi(y) \int_0^{\infty} \frac{e^{-\frac{(x-y)^2}{2\sigma^2}}}{\sqrt{2\pi}\sigma} dx \cdot (2h(y) - 1) dy \right) + c \\
 761 &= \left(\int_0^{\infty} \left(\int_{-\infty}^{\infty} \Pi(y) \frac{e^{-\frac{(x-y)^2}{2\sigma^2}}}{\sqrt{2\pi}\sigma} \cdot (2h(y) - 1) dy \right) dx \right) + c \\
 762 &= \left(\int_0^{\infty} \Pi'(x) \cdot (2H(x) - 1) dx \right) + c
 \end{aligned}$$

763 Replacing H with H' without loosing much in the approximation, we get that

$$765 \quad U_A(f) = \left(\int_0^{\infty} \frac{e^{-\frac{x^2}{2(\sigma^2+t^2)}}}{\sqrt{2\pi(\sigma^2+t^2)}} \cdot \frac{t^2}{\sigma^2+t^2} \frac{x}{d} dx \right) + c = \frac{t^2}{\sqrt{2\pi(\sigma^2+t^2)} \cdot d} + c$$

767 Similarly for $U_B(f)$ and hence, $U(f) = U_B(f) = U_A(f) = \frac{t^2}{\sqrt{2\pi(\sigma^2+t^2)} \cdot d} + c$.

768 E Proof of Theorem 8

769 We retain the instantiations of $h_A, h_B, \pi_A, \pi_B, c_A, c_B$ and σ as above. As seen above, in Scenario 2, $1_{x>0}$ is the classifier that optimizes Jury's utility and hence, $U(f_{\sigma}^*) = \frac{t^2}{\sqrt{2\pi(\sigma^2+t^2)} \cdot d} + c$.

770 Actually, it's approximately equal to this but the error is extremely small ($e^{-\Omega(d)}$, $d >> t, \sigma$).

771 In Scenario 1, the utility of any threshold classifier (f) with τ as the threshold is given by Equation 7 (without loss of generality, we can optimize over threshold classifiers). Therefore,

$$774 \quad U(f) = s_A \cdot \frac{t}{\sqrt{2\pi}d} e^{-(\tau - \sqrt{2\pi}\sigma_A)^2/2t^2} + s_B \cdot \frac{t}{\sqrt{2\pi}d} e^{-(\tau - \sqrt{2\pi}\sigma_B)^2/2t^2} + c$$

775 When $s_A = s_B = \frac{1}{2}$ and we assume that $|\sigma_A - \sigma_B| \leq \frac{t}{\sqrt{2\pi}}$, it's easy enough to see that the above expression is maximized at $\tau = \frac{\sqrt{2\pi}\sigma_A + \sqrt{2\pi}\sigma_B}{2}$. Therefore, the optimal classification accuracy in Scenario 1, is

$$776 \quad U(f_0^*) = \frac{t}{\sqrt{2\pi}d} e^{-(\frac{\sqrt{2\pi}\sigma_A - \sqrt{2\pi}\sigma_B}{2})^2/2t^2} + c$$

777 For $\sigma_B = \sigma$, $\sigma_A = 0.1\sigma$, $t = 0.9\sqrt{2\pi}\sigma$, $U(f_{\sigma}^*) > U(f_0^*)$.