Analysis of RPO

Quantitative Analysis of Simplified RPO

In this section we analyze a simplified version of the randomized positive override. Instead of the Bayes factor audit test that is based on the *ratios* of selection rates, we consider a model where a violation occurs if the absolute number of selected candidates from the majority class exceeds the number of selected candidates from the minority class by some threshold. Second, the probability of deviating from UNFAIR (which we call the *override* probability) is a fixed constant called p_{over}. Finally, rather than always flipping the coin to consider overriding each candidate from the group with lower acceptance rate (which we call the *protected* candidates), we only do so if the current difference in selection numbers exceeds some threshold, τ_0.

Algorithm 2 Randomized Positive Override (Simplified)

Input: Prediction C_t, utility U_t, violation V_t, class P_t
Output: New prediction D_t
if $V_t(\texttt{keep}) \geq \tau$ **then**
  $D_t = 0^{C_t}$
else if $P_t = 1 \wedge C_t = 0 \wedge V_t \geq \tau_0$ **then**
  $D_t \sim Bern\left(p_{\text{over}}\right)$
else
  $D_t = C_t$
end if

We can model this scenario as a random walk on the integers, where W_n, the position of the walker at time n, is the number of selected non-protected candidates minus the number of selected protected candidates. Moving to the right represents selecting a non-protected candidate, moving to the left represents selecting a protected candidate, and staying stationary represents rejecting a candidate (from either class).

The effect of the failsafe is represented as a *reflecting barrier* at the violation threshold b. The effect of the RPO is modeled by skewing the transition probabilities of the walker whenever $a \leq W_n < b$, where a is the threshold at which overrides begin. Without loss of generality, we can take $a = 0$, since we can always imagine shifting W_n, a, and b by some constant offset.

Concretely, in addition to b, transitions in the model are specified by the following parameters:

- p_f, q_f – transition probabilities at the failsafe:
$$P(W_{n+1} = W_n \,|\, W_n = b) = p_f$$
$$P(W_{n+1} = W_n - 1 \,|\, W_n = b) = q_f = 1 - p_f$$
- p, q – transition probabilities in RPO region:
$$P(W_{n+1} = W_n + 1 \,|\, 0 \leq W_n < b) = p$$
$$P(W_{n+1} = W_n - 1 \,|\, 0 \leq W_n < b) = q$$
$$P(W_{n+1} = W_n \,|\, 0 \leq W_n < b) = 1 - p - q$$
- p_c, q_c – base transition probabilities :
$$P(W_{n+1} = W_n + 1 \,|\, W_n < 0) = p_c$$
$$P(W_{n+1} = W_n - 1 \,|\, W_n < 0) = q_c$$
$$P(W_{n+1} = W_n \,|\, W_n < 0) = 1 - p_c - q_c$$

Note that these transition probabilities are determined by a number of underlying parameters of the full model. For example, the q transition probability is determined by the rate at which candidates of the two classes occur, the underlying classifier's probability of selecting a random protected candidate, and the override probability p_{over}.

We assume that $p_c > q_c$, meaning that without enforcement, the underlying classifier will tend toward violation. Additionally, we assume that $p_f \neq q_f$ and $p \neq q$, and that $q_f > 0$ and $q > 0$. An important quantity is the ratio $r = \frac{p}{q}$.

In addition, we have parameters for expected utility earned for each transition made in the various regions.

- μ – expected utility earned from a transition made in the RPO region
- μ_f – expected utility from a transition at the failsafe

Without loss of generality, we assume that the utility is 0 when transitioning outside of the RPO or failsafe regions. (Effectively, we measure the difference between utility with and without enforcement). Again, these average utilities are dictated by other parameters of the full model.

Let $Q(t)$ measure the utility earned at time t.

Theorem 2. *Under the RPO policy,*

$$\lim_{t\to\infty} \frac{\mathbb{E}\left[Q(t)\right]}{t} = \frac{\mu T + \mu_f F}{T + F + C}$$

where

$$T = \frac{-1}{p-q}\left(1 - \frac{\beta_0}{1-\beta_1}\right)$$
$$F = \frac{\beta_0}{q_f(1-\beta_1)}$$
$$C = \frac{1}{p_c - q_c}$$
$$\beta_0 = \frac{r^{b+1} - r^b}{r^{b+1} - 1}$$
$$\beta_1 = \frac{r^{b+1} - r}{r^{b+1} - 1}$$

The fraction $\frac{T}{T+F+C}$ measures the long run ratio of time spent in the RPO region, while $\frac{F}{T+F+C}$ is the ratio of time at the failsafe.

The rest of this section proves this theorem. We need a number of auxiliary lemmas first.

Lemma 3. *Let $W_t = i$ for some time t, where $0 \leq i < b$. Let E be the first time after t such that $W_E = -1$ or $W_E = b$. Then $P(W_E = b) = \gamma_i$, where*

$$\gamma_i = \frac{r^{b+1} - r^{b-i}}{r^{b+1} - 1}$$

Proof. Set $S_n = W_{t+n} - W_t$ and $A_n = \left(\frac{q}{p}\right)^{S_n}$. Then A_n is a martingale and E is a stopping time. We have that $A_0 = 1$, so by the martingale stopping theorem, we have:

$$\mathbb{E}\left[A_{E-n}\right] = 1$$

If at E, we have $W_E = -1$, then $S_{E-n} = -i-1$, while if $W_E = b$, we have $S_{E-n} = b - i$, hence

$$1 = \mathbb{E}[A_{E-n}] = (1-\gamma_i)\left(\frac{q}{p}\right)^{(-i-1)} + \gamma_i \left(\frac{q}{p}\right)^{b-i}$$

Solving for γ_i, we have:

$$\gamma_i = \frac{1 - \left(\frac{q}{p}\right)^{-i-1}}{\left(\frac{q}{p}\right)^{b-i} - \left(\frac{q}{p}\right)^{-i-1}}$$

Recalling that $r = \frac{p}{q}$, and rewriting in these terms, we have:

$$\gamma_i = \frac{r^{b+1} - r^{b-i}}{r^{b+1} - 1}$$

□

Note that $\gamma_0 = \beta_0$ and $\gamma_{b-1} = \beta_1$.

Lemma 4. *For $0 \le i, j \le b$, let s_{ij} be the expected number of times that the random walk, starting from state i will be in state j before reaching -1. Then*

$$s_{bb} = \frac{1}{q_f(1-\gamma_{b-1})}$$
$$s_{0b} = \frac{\gamma_0}{q_f(1-\gamma_{b-1})} = F$$

Proof. By the law of total expectation, conditioning on the first transition the walk makes from b, we have:

$$s_{bb} = 1 + p_f s_{bb} + q_f s_{(b-1)b}$$

For $s_{(b-1)b}$, we can condition on whether the walk ever hits b again before hitting -1 or not:

$$s_{(b-1)b} = \gamma_{b-1} s_{bb} + (1-\gamma_{b-1}) \cdot 0$$

By substituting this into the equation for s_{bb} and solving we obtain the stated result for s_{bb}. Finding s_{0b} is similar. □

Lemma 5. $\sum_{i=0}^{b-1} s_{0i} = T$*, where T is defined as in the statement of 2 In other words, if the walker is at state 0, then the expected number of transitions made in RPO states before next hitting -1 is T.*

Proof. WLOG assume the walker is at state 0 at time 0. Let S_n be the position of the walker at time n, let T_n be the number of transitions from states $0, \dots, b-1$ up to time n, and let F_n be the number of transitions made from state b up to time n. Set $X_n = S_n - T_n(p-q) - F_n(-q_f)$. Then $X_0 = 0$, and $X_0, X_1, \dots$ is a martingale with respect to $S_1, S_2, \dots$. Let N be the first time that the walker reaches state -1. Then N is a stopping time with respect to $S_1, S_2, \dots$, so that by the stopping time theorem, $\mathbb{E}[X_N] = 0$. But we also know that

$$\mathbb{E}[X_N] = \mathbb{E}[S_N] - \mathbb{E}[T_N](p-q) - \mathbb{E}[F_N](-q_f) = -1 - \left(\sum_{i=0}^{b-1} s_{0i}\right)(p-q) - s_{0b}(-q_f)$$

Setting the above expression equal to 0, and plugging in the value of s_{0b} in the previous lemma, we can solve for $\sum_{i=0}^{b-1} s_{0i}$ and find that it is $\frac{-1}{p-q}\left(1 - \frac{\beta_0}{1-\beta_1}\right) = T$. □

Lemma 6. *Let c to be the expected utility incurred for all transitions starting from when the walker is at 0 until it first hits -1 again. Then $c = \mu T + \mu_f F$.*

Proof. From the previous lemmas, we have seen that T will be the expected number of transitions made in RPO stages, and F will be the expected number of transitions made at the failsafe barrier. Since in expectation each transition from the former has utility μ, and each transition at the latter has utility μ_f, the total expected utility follows from Wald's equation. □

Lemma 7. *If the walker is at position -1, the expected number of transitions until it reaches 0 again is $\frac{1}{p_c - q_c} = C$.*

Proof. Similar to 4, by conditioning on the first transition. □

With all of these lemmas in place, we are now ready to prove 2.

Proof. (2)

Let B_0 be the first time at which the walker hits 0. Let N_0 be the time it next hits -1. Then define B_{i+1} to be the first time after N_i at which the walker hits 0, and define N_{i+1} to be the first time after B_{i+1} at which the walker hits -1.

Then the $V_{i+1} = N_{i+1} - N_i$ is an IID sequence, and

$$\mathbb{E}[V_{i+1}] = \mathbb{E}[N_{i+1} - N_i] = \mathbb{E}[N_{i+1} - B_{i+1}] - \mathbb{E}[B_{i+1} - N_i] = T + F + C$$

Let R_{i+1} be the utility earned from N_i to N_{i+1}. Then (V_i, R_i) form a renewal sequence with rewards, so that if $Q(t)$ is the total utility at time t,

$$\lim_{t\to\infty} \frac{\mathbb{E}[Q(t)]}{t} = \frac{\mathbb{E}[R_1]}{\mathbb{E}[V_1]} = \frac{\mu T + \mu F}{T + F + C}$$

□

When GBF is Optimal

In our experiments on utility maximization, we observed that GBF was frequently competitive with OPT. We conduct another set of experiments, with the goal to understand the conditions under which GBF maximizes utility. To answer this question, we synthesize and experiment with 700 datasets.

Datasets Each dataset belongs to one of three families: Beta, Pareto (power-law), or Exponential. Within a family, the scores for candidates coming from `C1` and `C2` are drawn from two, independently parameterized versions of the same distribution (e.g., within the Beta family, the scores of candidates from `C1` and the scores of candidates from `C2` come from two, independent Beta distributions, respectively). Letting μ_1 and μ_2 be the mean parameters of the Beta distributions for `C1` and `C2`, and letting $M = \{0.1, 0.3, 0.5, 0.7, 0.9\}$, we consider each pair $(\mu_1, \mu_2) \in M \times M$, i.e., each pair in the Cartesian product of the set M with itself. For each pair of mean parameters, we experiment with 3 different values

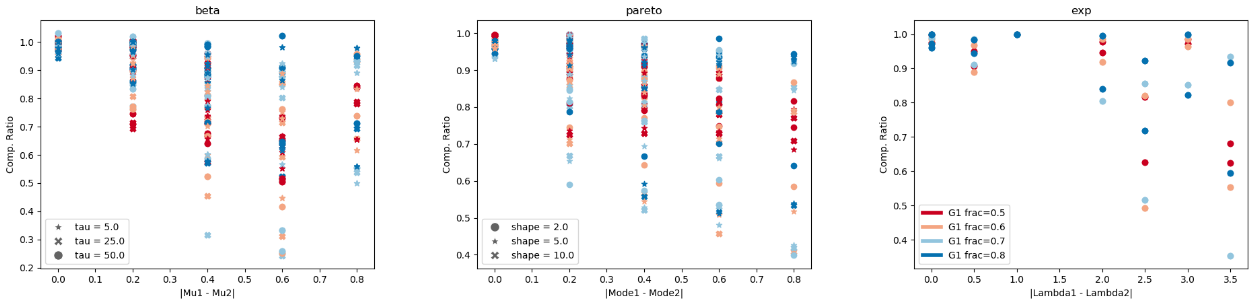

(a) Beta (b) Pareto (c) Exponential

Figure 4: **Competitive Ratios.** Each point in each plot represents the utility achieved by GBF divided by the utility achieved by OPT averaged over 10 shuffles of a dataset. The x-axis represents the absolute value of the difference between mean parameters of the score distributions for `G1` and `G2`. Color represents the fraction of the candidates from `G1` and mark type represents the value of the variance parameter (either τ or α), when applicable.

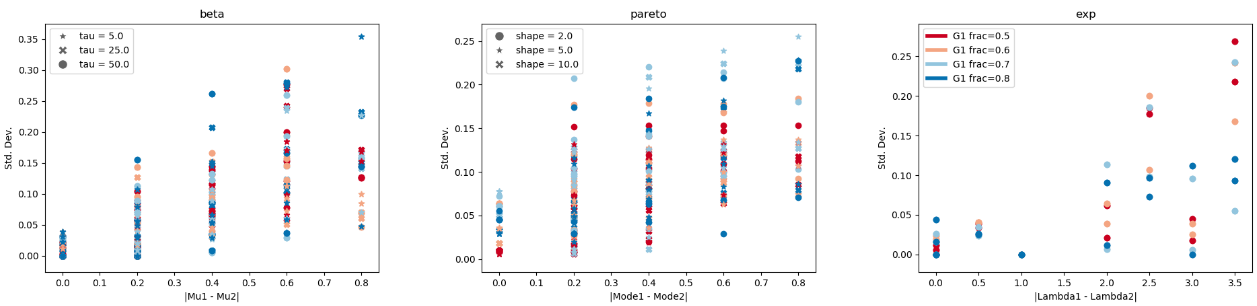

(a) Beta (b) Pareto (c) Exponential

Figure 5: **Standard Deviations.** Each point in each plot represents the standard deviation of the competivite ratio between GBF and OPT, over 10 shuffles of a dataset. As in Figure 4, the x-axis represents the absolute value of the difference between mean parameters of the score distributions for `G1` and `G2`. Color represents the fraction of the candidates from `G1` and mark type represents the value of the variance parameter (either τ or α), when applicable.

of τ^2: 5, 25, and 50 (τ is always the same for `C1` and `C2`). Similarly, for the Pareto family, letting m_1 and m_2 be the mode parameters of the Pareto distributions for `C1` and `C2`, respectively, we consider each pair $(m_1, m_2) \in M \times M$; we experiment with 3 different values of α^3: 2, 5, and 10 (α is always the same for `C1` and `C2`). Finally, for the Exponential family of datasets, let λ_1 and λ_2 be the rate parameters of `C1` and `C2`, respectively, and let $\Lambda = \{0.5, 1.0, 2.0, 3.0, 4.0\}$. Again, we consider each pair $(\lambda_1, \lambda_2) \in \Lambda \times \Lambda$. For each combination of parameters (e.g., (μ_1, μ_2, τ)), we generate 4 datasets, which differ in the fraction of candidates drawn from `C1`, either: 0.5, 0.6, 0.7, or 0.8. Overall, this yields 700 synthetic datasets[4]. We choose the Beta distributions because it can be used to model classifier confidences; the Pareto distribution because it arises naturally and has been studied in the context of fair selections (Raghavan et al. 2019); the exponential distribution to add additional variety to our experiments.

We run GBF and OPT on 10 random permutations of each dataset. For the Beta family, we set the score threshold, $\delta = 0.5$, and for the Pareto and Exponential families, δ is equal to the expected mean score of all candidates (respecting the proportion of `C1`). We use our first definition of utility (i.e., $U_t(\texttt{keep}) = |s_t - \delta|$). For each of the 7000 permutations, we calculate a competitive ratio by dividing the cumulative utility achieved by GBF by the cumulative utility achieved by OPT. Figure 4 visualizes the average competitive ratios (over the 10 randomizations) as a function of the difference in mean (mode, or rate) parameter of the score distributions for `C1` and `C2`. Interestingly, the plot reveals that when the means are close—no matter the value of the mean, mode or rate parameter, the distribution family, the variance of the distribution, or the fraction of the candidates from

[2]The Beta distribution can be parameterized by a mean, μ and an inverse-variance, τ.

[3]The Pareto distribution can be parameterized by a mode, μ, and a shape α.

[4]For Beta (and Pareto): 25 combinations of mean (mode) parameters $\times$ 3 values of τ (α) $\times$ 4 different fractions of candidate from `C1`; for Exponential: 25 combinations of rate parameters $\times$ 4 different fractions of `C1` candidates.

group C1—the average competitive ratio for GBF is high. This suggests that when the average scores of the candidates from C1 and C2 are similar, GBF is likely to be sufficient in terms of maximizing utility. On the other hand, when there is significant disparity in the average scores of candidates from C1 and C2, there is more uncertainty in the competitiveness of GBF. That said, even when significant disparity exists, under some score distributions, GBF remains highly competitive with OPT. However, no global trends are observed when uncertainty (i.e., τ and α) or the fraction of candidates from C1 are increased/decreased. We note that when the difference between mean parameter is low, the standard deviation of the competitive ratio is also low; as the difference increases, so does the spread of standard deviations across experimental conditions (Figure 5).