Appendix A: Omitted Details from Section 2.2

A.1. Proof of Proposition 1

This proof adapts the proof of Theorem 7.1 in Mehta et al. (2007b) to our setting, which involves presenting a ranking to arriving volunteers as well as a more general concave objective function. Capacities in their setting can be viewed of as an opportunity's equity threshold in our setting, and the only difference lies in the reward of γ for a sign-up (i.e., a match) beyond an opportunity's equity threshold (i.e., capacity). To show that no online algorithm obtains a competitive ratio greater than $\frac{e-1+\gamma}{e}$, it is sufficient to show that there exists a distribution over a class of instances for which no deterministic online algorithm obtains more than a $\frac{e-1+\gamma}{e}$ fraction of the optimal clairvoyant benchmark (Yao 1977).

For a fixed efficiency weight γ, consider an instance with n opportunities indexed $i \in \{1,\dots,n\}$ with homogeneous equity thresholds of m. In this instance there are exactly $n \cdot m$ arriving volunteers, each with the same choice model: an arriving volunteer only considers the top-ranked opportunity and will sign up if and only if it is *compatible* with that opportunity. The compatibility structure is defined as follows: the first m volunteers are compatible with all opportunities, the next m volunteers are only compatible with opportunities $i > 1$, and so on. In other words, volunteers $(j-1)\cdot m+1$ through $j \cdot m$ are only compatible with opportunities $i \geqslant j$.

We now consider the class of instances that we can obtain by permuting the indices of the opportunities in the instance described above. Specifically, we apply a permutation ρ to the indices of opportunities, such that they are now indexed $\{\rho(1),\dots,\rho(n)\}$. Given this indexing, volunteers $(j-1)\cdot m+1$ through $j \cdot m$ are only compatible with opportunities i where $\rho^{-1}(i) \geqslant j$.

In any such permuted instance, the optimal clairvoyant benchmark obtains a reward of $n \cdot m$. To see this, consider a clairvoyant solution that for each volunteer $t \in \{(j-1)\cdot m+1,\dots,j\cdot m\}$, it presents a ranking with opportunity $\rho^{-1}(j)$ ranked first and the remaining opportunities ranked arbitrarily. These m volunteers are all compatible with the top-ranked opportunity, and thus they deterministically sign up. Each opportunity then gets exactly m sign-ups, leading to an objective value of $n \cdot m$.

We now bound the performance of any deterministic online algorithm on the uniform distribution over this class of instances.

Claim 1 *Suppose that the permutation ρ is drawn uniformly at random from the set of all permutations of n opportunities. For any deterministic online algorithm π,*

$$\mathbb{E}[\pi] \quad \leqslant \quad \left(\frac{e-1+\gamma}{e}\right) n \cdot m, \tag{13}$$

where the expectation is taken over permutations ρ.

Proof of Claim 1: To aid in this proof, let us define $v_{i,j}$ as the number of volunteers from the set $\{(j-1)\cdot m+1,\dots,j\cdot m\}$ that sign up for opportunity $\rho^{-1}(i)$. Based on the compatibility structure, no online algorithm can distinguish between opportunities that are still compatible with arriving volunteers. Therefore, we have:

$$E_\rho[v_{i,j}] \quad \leqslant \quad \begin{cases} \dfrac{m}{n-j+1}, & \text{if } i \geqslant j \\ 0 & \text{if } i < j \end{cases}$$

To see this, note that for each volunteer t between $t \in \{(j-1)\cdot m+1, j\cdot m\}$, there are a total of $n-j+1$ compatible opportunities, i.e., any opportunity i where $\rho^{-1}(i) \geqslant j$. Since the permutation is chosen uniformly at random, all remaining compatible opportunities appear identical to the online algorithm. This means that the number of volunteer sign-ups for each opportunity cannot exceed the average of $\frac{m}{n-j+1}$, when taking an expectation over permutations.

After all volunteers arrive, opportunity $\rho^{-1}(i)$ has received at most $\sum_{j=1}^{i} \frac{m}{n-j+1}$ sign-ups. As the objective function is concave in the number of sign-ups for each opportunity, the expected objective value for an online algorithm π is at most:

$$E_\rho[\pi] \quad \leqslant \quad \sum_{i=1}^{n} \min\left\{m, \sum_{j=1}^{i} \frac{m}{n-j+1}\right\} + \gamma \max\left\{0, \sum_{j=1}^{i} \frac{m}{n-j+1} - m\right\} \quad := \quad \text{RHS} \tag{14}$$

Taking the limit as n and m approach infinity, we have

$$\lim_{n,m\to\infty} \frac{\text{RHS}}{n\cdot m} \quad = \quad \left(\frac{e-1+\gamma}{e}\right).$$

This completes the proof of Claim 1. □

In combination with Yao's lemma, this proves that no online algorithm (deterministic or randomized) can obtain a competitive ratio greater than $\frac{e-1+\gamma}{e}$. To see that this holds for any finite minimum equity threshold $\underline{m}$, note that we can simply add an additional opportunity with an arbitrary equity threshold and no compatible volunteers to the set of instances described. □

A.2. Proof of Theorem 1

We will prove this result in three steps. (i) First, we will define pseudo-rewards that depend on the particular choice of penalty function used by a balancing algorithm π^{BAL} (see Definition 3). (ii) Then, in Lemma 1 we will show that the expected value of the algorithm is at least a constant fraction of the expected sum of the pseudo-rewards. (iii) We subsequently show in Lemma 2 that the expected sum of the pseudo-rewards is at least a constant fraction of the expected value of the clairvoyant benchmark (see Definition 1). Combining the three steps establishes a lower bound on the competitive ratio of π^{BAL} (see Definition 2). The bound on the competitive ratio is equal to the product of the constant-factors from Steps (ii) and (iii), which are both parameterized by γ and $\underline{m}$.

Step (i): Defining Pseudo-Rewards. Consider a balancing algorithm π^{BAL} with penalty function ψ, which is a non-increasing function of the "fill rate" of an opportunity, i.e., the ratio between its number of sign-ups and its equity threshold. For such an algorithm, we will define pseudo-rewards which we emphasize are purely for accounting purposes. In the subsequent steps of the proof, we will use the pseudo-rewards as an intermediary to compare the performance of π^{BAL} with that of the clairvoyant benchmark.

Recall that $X_{i,t}(\vec{S})$ is an indicator random variable that equals 1 if and only if volunteer t signs up for opportunity i when presented with ranking $\vec{S}$. For each arriving volunteer $t \in \{1,\ldots,T\}$, let $\vec{S}_t^{\text{BAL}}$ denote the ranking provided by π^{BAL}, and let $\vec{S}_t^{\text{OPT}}$ denote the ranking provided by our clairvoyant benchmark. Furthermore, recall that $Z_{i,t}$ denotes the number of sign-ups for opportunity i after the arrival of volunteer

t, and let $f_{i,t} = Z_{i,t}/m_i$ denote the fill rate of the opportunity after the arrival of volunteer t. With this notation in mind, we define

$$\lambda_t = \sum_{i=1}^{n} \psi(f_{i,t-1}) X_{i,t}(\vec{S}_t^{\mathrm{OPT}}) \tag{15}$$

$$\theta_i = \sum_{t=1}^{T} \Big(1-\psi(f_{i,t-1})\Big) \mathbb{1}\{Z_{i,t} \leqslant m_i\} X_{i,t}(\vec{S}_t^{\mathrm{BAL}}) \tag{16}$$

We note that these pseudo-rewards are random and depend on the volunteers' choices.

Step (ii): Lower Bounding $\mathbb{E}[\pi^{\mathbf{BAL}}(\mathcal{I})]$ with Pseudo-Rewards. We use the pseudo-rewards defined in Step (i) to establish a lower bound on the expected value of π^{BAL}, as formalized in the following lemma.

Lemma 1 *Consider a balancing algorithm π^{BAL} with penalty function ψ. For any instance $\mathcal{I}$ with an efficiency weight γ, we have*

$$\mathbb{E}[\pi^{\mathrm{BAL}}(\mathcal{I})] \geqslant \alpha \cdot \mathbb{E}\Big[\sum_{i=1}^{n} \theta_i + \sum_{t=1}^{T} \lambda_t\Big],$$

where $\alpha = \min\{1, \gamma/\psi(1)\}$ and where the expectation is taken over the volunteers' choices.[13]

Proof of Lemma 1 Recall that (OBJ) can be re-written as the objective function in Equation (1). Taking the expectation over the volunteers' choices, we have the following series of inequalities:

$$\mathbb{E}[\pi^{\mathrm{BAL}}(\mathcal{I})] = \mathbb{E}\Big[\sum_{i=1}^{n}\sum_{t=1}^{T} \mathbb{1}\{Z_{i,t} \leqslant m_i\} X_{i,t}(\vec{S}_t^{\mathrm{BAL}}) + \gamma \mathbb{1}\{Z_{i,t} > m_i\} X_{i,t}(\vec{S}_t^{\mathrm{BAL}})\Big] \tag{17}$$

$$= \mathbb{E}\Big[\sum_{i=1}^{n}\sum_{t=1}^{T} \Big(1-\psi(f_{i,t-1})\Big) \mathbb{1}\{Z_{i,t} \leqslant m_i\} X_{i,t}(\vec{S}_t^{\mathrm{BAL}}) + \psi(f_{i,t-1}) \mathbb{1}\{Z_{i,t} \leqslant m_i\} X_{i,t}(\vec{S}_t^{\mathrm{BAL}}) + \gamma \mathbb{1}\{Z_{i,t} > m_i\} X_{i,t}(\vec{S}_t^{\mathrm{BAL}})\Big] \tag{18}$$

$$= \mathbb{E}\Big[\sum_{i=1}^{n} \theta_i + \sum_{i=1}^{n}\sum_{t=1}^{T} \psi(f_{i,t-1}) \mathbb{1}\{Z_{i,t} \leqslant m_i\} X_{i,t}(\vec{S}_t^{\mathrm{BAL}}) + \gamma \mathbb{1}\{Z_{i,t} > m_i\} X_{i,t}(\vec{S}_t^{\mathrm{BAL}})\Big] \tag{19}$$

$$\geqslant \mathbb{E}\Big[\sum_{i=1}^{n} \theta_i + \sum_{i=1}^{n}\sum_{t=1}^{T} \psi(f_{i,t-1}) \mathbb{1}\{Z_{i,t} \leqslant m_i\} X_{i,t}(\vec{S}_t^{\mathrm{BAL}}) + \alpha \cdot \psi(f_{i,t-1}) \mathbb{1}\{Z_{i,t} > m_i\} X_{i,t}(\vec{S}_t^{\mathrm{BAL}})\Big] \tag{20}$$

$$\geqslant \alpha \cdot \mathbb{E}\Big[\sum_{i=1}^{n} \theta_i + \sum_{i=1}^{n}\sum_{t=1}^{T} \psi(f_{i,t-1}) \mathbb{1}\{Z_{i,t} \leqslant m_i\} X_{i,t}(\vec{S}_t^{\mathrm{BAL}}) + \psi(f_{i,t-1}) \mathbb{1}\{Z_{i,t} > m_i\} X_{i,t}(\vec{S}_t^{\mathrm{BAL}})\Big] \tag{21}$$

$$= \alpha \cdot \mathbb{E}\Big[\sum_{i=1}^{n} \theta_i + \sum_{i=1}^{n}\sum_{t=1}^{T} \psi(f_{i,t-1}) X_{i,t}(\vec{S}_t^{\mathrm{BAL}})\Big] \tag{22}$$

$$\geqslant \alpha \cdot \mathbb{E}\Big[\sum_{i=1}^{n} \theta_i + \sum_{i=1}^{n}\sum_{t=1}^{T} \psi(f_{i,t-1}) X_{i,t}(\vec{S}_t^{\mathrm{OPT}})\Big] \tag{23}$$

$$= \alpha \cdot \mathbb{E}\Big[\sum_{i=1}^{n} \theta_i + \sum_{t=1}^{T} \lambda_t\Big] \tag{24}$$

There are two key steps in this series of inequalities. First, in Line (19), we rely on the fact that whenever the indicator $\mathbb{1}\{Z_{i,t} > m_i\}$ is 1, we must have $f_{i,t-1} \geqslant 1$. Furthermore,

$$f_{i,t-1} \geqslant 1 \quad \Longrightarrow \quad \alpha \cdot \psi(f_{i,t-1}) \leqslant \alpha \cdot \psi(1) = \min\{\psi(1), \gamma\} \leqslant \gamma,$$

[13] If the denominator $\psi(1) = 0$, we set $\alpha = 1$.

where the first inequality holds because ψ is weakly decreasing.

Then, in Line (23), we use the optimality criteria of the exponential balancing function π^{BAL}, namely, that π^{BAL} presents a ranking that maximizes $\sum_{i\in\vec{S}}\psi(Z_{i,t})\cdot\phi_{i,t}(\vec{S})$. This implies that

$$\mathbb{E}\Big[\sum_{i=1}^{n}\sum_{t=1}^{T}\psi(f_{i,t-1})X_{i,t}(\vec{S}_t^{\mathrm{BAL}})\Big] \quad\geqslant\quad \mathbb{E}\Big[\sum_{i=1}^{n}\sum_{t=1}^{T}\psi(f_{i,t-1})X_{i,t}(\vec{S}_t^{\mathrm{OPT}})\Big],$$

due to the fact that $\mathbb{E}[X_{i,t}(\vec{S})]=\phi_{i,t}(\vec{S})$ and the tower property of expectations. (These steps are identical to the rigorous application of the tower property of expectations in a similar proof in Manshadi et al. (2022b).)

□

Step (iii): Upper Bounding $\mathbb{E}[\mathbf{OPT}(\mathcal{I})]$ with Pseudo-Rewards. We now establish an upper bound on the clairvoyant solution for any instance using the expected sum of the pseudo-rewards under π^{BAL}.

Lemma 2 *Consider a balancing algorithm π^{BAL} with penalty function ψ. For any instance $\mathcal{I}$ with an efficiency weight γ and a minimum equity threshold $\underline{m}$, we have*

$$\mathbb{E}\Big[\sum_{i=1}^{n}\theta_i+\sum_{t=1}^{T}\lambda_t\Big] \quad\geqslant \beta\cdot\kappa\cdot\mathbb{E}[\mathrm{OPT}(\mathcal{I})], \tag{25}$$

where the expectation is taken over the volunteers' choices, $\beta=\min\Big\{1,\lim_{x\to\infty}\frac{\psi(x)}{\gamma\cdot\psi(0)}\Big\}$,[14] *and where*

$$\kappa=\min_{m,Z\in\mathbb{Z}_+,m\geqslant\underline{m}}\psi(Z/m)+(1/m)\sum_{k=1}^{\min\{Z,m\}}\Big(1-\psi\big((k-1)/m\big)\Big). \tag{26}$$

Proof of Lemma 2 We will show that a tighter version of Equation (25) holds for each opportunity and for any realization of volunteer choices (similar to the approach taken in, e.g., Manshadi et al. 2022b). Specifically, letting $\mathcal{O}_i$ denote the set of volunteers who sign up for opportunity i under the optimal clairvoyant policy, we will show

$$\theta_i+\sum_{t\in\mathcal{O}_i}\lambda_t\geqslant\beta\cdot\kappa\cdot\Big(\gamma|\mathcal{O}_i|+(1-\gamma)\min\{m_i,|\mathcal{O}_i|\}\Big) \tag{27}$$

Summing over all opportunities and taking expectation over the volunteers' choices completes the proof.

To establish Equation (27), we will separately bound the value of θ_i and the value of the λ_t pseudo-rewards, for all $t\in\mathcal{O}_i$ (i.e., for all volunteers that sign up for opportunity i under the clairvoyant solution). Starting with the definition of λ_t in Equation (15),

$$\begin{aligned}\sum_{t\in\mathcal{O}_i}\lambda_t &= \sum_{t\in\mathcal{O}_i}\psi(f_{i,t-1})X_{i,t}(\vec{S}_t^{\mathrm{OPT}}) &(28)\\ &= \sum_{t\in\mathcal{O}_i}\psi(f_{i,t-1}) &(29)\\ &\geqslant \sum_{t\in\mathcal{O}_i}\psi(f_{i,T}) &(30)\\ &= |\mathcal{O}_i|\psi(f_{i,T}) &(31)\end{aligned}$$

Equality in Line (29) follows from the definition that $t\in\mathcal{O}_i$ only if the volunteer signs up for opportunity i when presented with the optimal clairvoyant benchmark's ranking, i.e., $X_{i,t}(\vec{S}_t^{\mathrm{OPT}})=1$. The following

[14] If the denominator $\gamma\cdot\psi(0)=0$, we set $\beta=1$.

inequality in Line (30) relies on the fact that ψ is a weakly decreasing function and $f_{i,T} \geqslant f_{i,t-1}$ for all $t \in \{1,\dots,T\}$.

Next, based on the definition of θ_i in Equation (16),

$$\theta_i = \sum_{t=1}^{T} \Big(1-\psi(f_{i,t-1})\Big)\mathbb{1}\{Z_{i,t} \leqslant m_i\} X_{i,t}(\vec{S}_t^{\text{BAL}}) \tag{32}$$

$$= \sum_{k=1}^{Z_{i,T}} \Big(1-\psi\big((k-1)/m_i\big)\Big)\mathbb{1}\{k \leqslant m_i\} \tag{33}$$

$$= \sum_{k=1}^{\min\{Z_{i,T},m_i\}} \Big(1-\psi\big((k-1)/m_i\big)\Big) \tag{34}$$

$$= \sum_{k=1}^{m_i\cdot\min\{f_{i,T},1\}} \Big(1-\psi\big((k-1)/m_i\big)\Big) \tag{35}$$

Based on our separate bounds for the pseudo-rewards, to complete the proof of Lemma 2, it is sufficient to establish that for all $|\mathcal{O}_i| \in \mathbb{Z}_+$ and all $f_{i,T} \in \mathbb{R}_+$, we have

$$|\mathcal{O}_i|\psi(f_{i,T}) + \sum_{k=1}^{m_i\cdot\min\{f_{i,T},1\}} \Big(1-\psi\big((k-1)/m_i\big)\Big) \quge \beta\cdot\kappa\cdot\Big(\gamma|\mathcal{O}_i| + (1-\gamma)\min\{m_i,|\mathcal{O}_i|\}\Big) \tag{36}$$

To prove Inequality (36), we define two function, $L_i(y)$ and $R_i(y)$ for $y \in \mathbb{R}_+$, which represent the left and right hand sides of the inequality as a continuous function of $|\mathcal{O}_i|$. Specifically,

$$L_i(y) = y\cdot\psi(f_{i,T}) + \sum_{k=1}^{m_i\cdot\min\{f_{i,T},1\}} \Big(1-\psi\big((k-1)/m_i\big)\Big) \tag{37}$$

$$R_i(y) = \beta\cdot\kappa\cdot\Big(\gamma\cdot y + (1-\gamma)\min\{m_i,y\}\Big) \tag{38}$$

We make the following structural observations about the two functions: $L_i(y)$ is a linear function with $L_i(0) \geqslant 0$ and slope $\psi(f_{i,T})$. Furthermore, $R_i(y)$ is a piecewise-linear function with $R_i(0) = 0$, with slope $\beta\cdot\kappa$ for $y \in [0, m_i]$ and slope $\beta\cdot\kappa\cdot\gamma$ for $y > m_i$. Given these observations, to prove Inequality (36), it is sufficient to show that (i) $L_i(m_i) \geqslant R_i(m_i)$, and (ii) for $y > m_i$, the slope of $R_i(y)$ does not exceed that of $L_i(y)$.

To establish (i) $L_i(m_i) \geqslant R_i(m_i)$, we have

$$m_i\psi(f_{i,T}) + \sum_{k=1}^{m_i\cdot\min\{f_{i,T},1\}} \Big(1-\psi\big((k-1)/m_i\big)\Big)$$

$$= m_i\left(\psi(f_{i,T}) + (1/m_i)\sum_{k=1}^{m_i\cdot\min\{f_{i,T},1\}} \Big(1-\psi\big((k-1)/m_i\big)\Big)\right) \tag{39}$$

$$\geqslant m_i\left(\min_{m,z\in\mathbb{Z}_+,m\geqslant\underline{m}} \psi(z/m) + (1/m)\sum_{k=1}^{\min\{z,m\}} \Big(1-\psi\big((k-1)/m\big)\Big)\right) \tag{40}$$

$$= m_i\cdot\kappa \tag{41}$$

$$\geqslant m_i\cdot\beta\cdot\kappa \tag{42}$$

To establish (ii) for $y > m_i$, the slope of $R_i(y)$ does not exceed that of $L_i(y)$, we have

$$\beta\cdot\kappa\cdot\gamma \leqslant \beta\cdot\psi(0)\cdot\gamma = \min\{\psi(0)\cdot\gamma, \lim_{x\to\infty}\psi(x)\} \leqslant \psi(f_{i,T}),$$

where the first inequality holds by plugging in $y = 0$ to the definition if κ in Equation 26 and the last inequality holds because ψ is weakly decreasing.

Therefore, we have shown that Inequality (36) holds for all $|\mathcal{O}_i| \in \mathbb{Z}_+$ and all $f_{i,T} \in \mathbb{R}_+$, which is sufficient to complete the proof of Lemma 2. □

Based on the pseudo-rewards defined in Step (i) and using Lemmas 1 and 2, we can establish the following series of inequalities to bound the competitive ratio of a balancing algorithm π^{BAL} with penalty function ψ:

$$\mathbb{E}[\pi^{\text{BAL}}(\mathcal{I})] \quad \geqslant \quad \alpha \cdot \mathbb{E}\Big[\sum_{i=1}^{n} \theta_i + \sum_{t=1}^{T} \lambda_t\Big] \quad \geqslant \alpha \cdot \beta \cdot \kappa \cdot \mathbb{E}[\text{OPT}(\mathcal{I})], \tag{43}$$

where $\alpha = \min\left\{1, \frac{\gamma}{\psi(1)}\right\}$, $\beta = \min\left\{1, \lim_{x\to\infty} \frac{\psi(x)}{\gamma\cdot\psi(0)}\right\}$, and κ is defined in Equation (26). The first inequality comes from Lemma 1, and the second comes from Lemma 2. This completes the proof of Theorem 1. □

A.3. Proof of Corollary 1

By Theorem 1, the competitive ratio of a balancing algorithm consists of three terms. For the specific penalty function $\psi(x) = 1 - (1-\gamma)\exp(-1 + \min\{x, 1\})$, we have $\psi(0) = 1 - (1-\gamma)e^{-1}$ and $\psi(1) = \lim_{x\to\infty} \psi(x) = \gamma$. Therefore $\alpha = \min\left\{1, \frac{\gamma}{\psi(1)}\right\} = 1$ and $\beta = \min\left\{1, \lim_{x\to\infty} \frac{\psi(x)}{\gamma\cdot\psi(0)}\right\} = 1$.

For the final term,

$$\begin{aligned}
\kappa &= \min_{m,z\in\mathbb{Z}_+, m\geqslant \underline{m}} \psi(z/m) + (1/m) \sum_{k=1}^{\min\{z,m\}} \Big(1 - \psi\big((k-1)/m\big)\Big) && (44)\\
&\geqslant \min_{m,z\in\mathbb{Z}_+, m\geqslant \underline{m}} \psi(z/m) + (1/m)e^{-1/\underline{m}} \sum_{k=1}^{\min\{z,m\}} \Big(1 - \psi\big(k/m\big)\Big) && (45)\\
&\geqslant e^{-1/\underline{m}} \min_{m,z\in\mathbb{Z}_+, m\geqslant \underline{m}} \psi(z/m) + (1/m) \sum_{k=1}^{\min\{z,m\}} \Big(1 - \psi\big(k/m\big)\Big) && (46)\\
&\geqslant e^{-1/\underline{m}} \min_{m,z\in\mathbb{Z}_+, m\geqslant \underline{m}} \psi(z/m) + \int_0^{\min\{z/m,1\}} 1 - \psi(x)\ \partial x && (47)\\
&= e^{-1/\underline{m}} \min_{m,z\in\mathbb{Z}_+, m\geqslant \underline{m}} 1 - (1-\gamma)\, e^{-1+\min\{z/m,1\}} + (1-\gamma)\, e^{-1+\min\{z/m,1\}} - (1-\gamma)\, e^{-1} && (48)\\
&= e^{-1/\underline{m}} \min_{m,z\in\mathbb{Z}_+, m\geqslant \underline{m}} 1 - (1-\gamma)\, e^{-1} && (49)\\
&= e^{-1/\underline{m}} \left(\frac{e-1+\gamma}{e}\right) && (50)
\end{aligned}$$

The inequality in Line (45) relies on the algebraic fact that for this penalty function, $1 - \psi((k-1)/m_i) \geqslant e^{-1/\underline{m}}(1 - \psi(k/m_i))$ for all $k \in \mathbb{N}$. The inequality in Line (47) comes from the fact that $\psi(x)$ is weakly decreasing in x, and hence the prior summation is equivalent to a right Riemann sum of a weakly increasing function, which is then lower-bounded by the integral.

Due to the fact that $\lim_{\underline{m}\to\infty} e^{-1/\underline{m}} \left(\frac{e-1-\gamma}{e}\right) = \frac{e-1-\gamma}{e}$, this algorithm asymptotically obtains the best-possible competitive ratio, as it matches the upper bound for any online algorithm that was established in Proposition 1. □

A.4. Proof of Corollary 2

The proof of Corollary 2 is nearly identical to the proof of Theorem 1, replacing the ranking given by the *exact* balancing algorithm $\vec{S}_t^{\text{BAL}}$ with the ranking given by the *approximate* balancing algorithm $\vec{S}_t^{\text{BAL-}\zeta}$. The

only difference comes in Step (ii), Line (23), where we must introduce a multiplicative factor of ζ. Specifically, based on the approximate optimality of the ranking $\vec{S}_t^{\text{BAL-}\zeta}$, we have

$$\mathbb{E}\Big[\sum_{i=1}^{n}\sum_{t=1}^{T}\psi(f_{i,t-1})X_{i,t}(\vec{S}_t^{\text{BAL-}\zeta})\Big] \geqslant \zeta\cdot\mathbb{E}\Big[\sum_{i=1}^{n}\sum_{t=1}^{T}\psi(f_{i,t-1})X_{i,t}(\vec{S}_t^{\text{BAL}})\Big] \geqslant \zeta\cdot\mathbb{E}\Big[\sum_{i=1}^{n}\sum_{t=1}^{T}\psi(f_{i,t-1})X_{i,t}(\vec{S}_t^{\text{OPT}})\Big],$$

The remainder of the proof follows an identical path as the proof of Theorem 1, with the only difference being the inclusion of a multiplicative factor ζ. □

Appendix B: Omitted Details from Section 3.2

We begin in Appendix B.1 by formally proving Proposition 2. In Appendix B.2, we discuss the performance of similar score-based rankings under other choice models. Appendix B.3 investigates the impact of appeal misspecification, and Appendix B.4 contains a formal discussion around the impact of c_1. In Appendix B.5 we provide additional details about the simulations in Figure 4, and in Appendix B.6 we present additional sets of simulations. Lastly, Appendix B.7 contains a thought experiment that builds on the discussion in Step 4 of Section 3.2.

B.1. Proof of Proposition 2

We prove Proposition 2 in two steps. First, we will place a lower bound on the objective function of (PRO) for the score-based ranking $\vec{S}^{\text{SB}}$ (defined in Proposition 2). Then, we will place an upper bound on the maximum value of that objective function. Comparing the two bounds establishes the approximation guarantee stated in Proposition 2. As the proof focuses on a single representative volunteer t, we simplify notation and drop the subscript t. Similarly, we will replace $\psi(Z_{i,t-1}/m_i)$ with ψ_i.

To aid in this proof, we first note that in Definition 4, we have normalized the appeal of the no-choice option to 1 (without loss of generality). This implies that the probability of no-choice under ranking $\vec{S}$ is given by

$$\phi_0(\vec{S}) \quad = \quad 1-\sum_{i\in\vec{S}}\phi_i(\vec{S}) \quad = \quad \frac{1}{1+\sum_{j\in\vec{S}}\theta_{k(\vec{S},j)}\cdot a_j}. \tag{51}$$

Let us define $M := \max_{\vec{S}\in\mathcal{S}} 1+\sum_{j\in\vec{S}}\theta_{k(\vec{S},j)}\cdot a_j$. This is an upper bound on the denominator in Equation (51). Recall that $\bar{p}$ represents an upper bound on the probability of volunteer t signing up for any opportunity. We will leverage the following relationship between $\bar{p}$ and M:

$$\bar{p} \quad \geqslant \quad \max_{\vec{S}\in\mathcal{S}}\sum_{i\in\vec{S}}\phi_i(\vec{S}) \quad = \quad 1-\min_{\vec{S}\in\mathcal{S}}\phi_0(\vec{S}) \quad = \quad 1-\frac{1}{M} \tag{52}$$

Equivalently, $1\geqslant M(1-\bar{p})$.

Step 1: Lower Bound on Objective of (PRO) for $\vec{S}^{\text{SB}}$. Consider the ranking $\vec{S}^{\text{SB}}$ that ranks opportunities in descending order of $\psi_i\cdot a_i$. In that case, we have

$$\sum_{i\in\vec{S}^{\text{SB}}}\psi_i\cdot\phi_i(\vec{S}^{\text{SB}}) \quad = \quad \frac{\sum_{i\in\vec{S}^{\text{SB}}}\psi_i\cdot\theta_{k(\vec{S}^{\text{SB}},i)}\cdot a_i}{1+\sum_{i\in\vec{S}^{\text{SB}}}\cdot\theta_{k(\vec{S}^{\text{SB}},i)}\cdot a_i} \quad \geqslant \quad \frac{\sum_{i\in\vec{S}^{\text{SB}}}\psi_i\cdot\theta_{k(\vec{S}^{\text{SB}},i)}\cdot a_i}{M} \tag{53}$$

The equality comes from applying the definition of choice probability under the MNL model with position bias, given by Equation (6). The inequality comes from applying the definition of M, which is an upper bound on the denominator.

Step 2: Upper Bound on Optimal Solution to (PRO). Consider any ranking $\vec{S}$. We have:

$$\begin{aligned}
\sum_{i\in\vec{S}} \psi_i \cdot \phi_i(\vec{S}) &= \frac{\sum_{i\in\vec{S}} \psi_i \cdot \theta_{k(\vec{S},i)} \cdot a_i}{1+\sum_{i\in\vec{S}} \cdot \theta_{k(\vec{S},i)} \cdot a_i} \\
&\leqslant \frac{\sum_{i\in\vec{S}} \psi_i \cdot \theta_{k(\vec{S},i)} \cdot a_i}{1} \\
&\leqslant \frac{\sum_{i\in\vec{S}^{\text{SB}}} \psi_i \cdot \theta_{k(\vec{S}^{\text{SB}},i)} \cdot a_i}{1} \qquad (54) \\
&\leqslant \frac{\sum_{i\in\vec{S}^{\text{SB}}} \psi_i \cdot \theta_{k(\vec{S}^{\text{SB}},i)} \cdot a_i}{M(1-\bar{p})} \qquad (55)
\end{aligned}$$

Equality in the first line comes from applying the definition of choice probability under the MNL model with position bias, given by Equation (6). The inequality in Line (54) relies on Claim 2 (presented below), while the inequality in Line (55) comes from the definitions of M and $\bar{p}$, as described in Equation (52).

Combining Steps 1 and 2, we see that for any ranking $\vec{S}$,

$$\sum_{i\in\vec{S}^{\text{SB}}} \psi_i \cdot \phi_i(\vec{S}^{\text{SB}}) \geqslant \frac{\sum_{i\in\vec{S}^{\text{SB}}} \psi_i \cdot \theta_{k(\vec{S}^{\text{SB}},i)} \cdot a_i}{M} \geqslant (1-\bar{p}) \sum_{i\in\vec{S}} \psi_i \cdot \phi_i(\vec{S})$$

To complete the proof of Proposition 2, all that remains is to prove the following claim:

Claim 2 (Optimality Condition of Score-Based Ranking) *The score-based ranking $\vec{S}^{\text{SB}}$ which ranks opportunities in descending order*[15] *of $\psi_i \cdot a_i$ satisfies*

$$\vec{S}^{\text{SB}} \in \text{argmax}_{\vec{S}\in\mathcal{S}} \sum_{i\in\vec{S}} \psi_i \cdot \theta_{k(\vec{S},i)} \cdot a_i \qquad (56)$$

Proof of Claim 2 We show the optimality of this score-based ranking by a contradiction argument. Suppose there is a ranking $\vec{S} \neq \vec{S}^{\text{SB}}$ which obtains a strictly higher value for the summation in Equation (56). Since $\vec{S} \neq \vec{S}^{\text{SB}}$, there must exist opportunities i and j such that i has a better rank than opportunity j even though opportunity i has the worse score, i.e., $\theta_{k(\vec{S},i)} \geqslant \theta_{k(\vec{S},j)}$ but $\psi_i \cdot a_i \leqslant \psi_j \cdot a_j$. If we swap the ranks of opportunities i and j, the summation in Equation (56) will change by

$$\theta_{k(\vec{S},i)} \cdot \psi_j \cdot a_j - \theta_{k(\vec{S},j)} \cdot \psi_j \cdot a_j + \theta_{k(\vec{S},j)} \cdot \psi_i \cdot a_i - \theta_{k(\vec{S},i)} \cdot \psi_i \cdot a_i = (\theta_{k(\vec{S},i)} - \theta_{k(\vec{S},j)})(\psi_j \cdot a_j - \psi_i \cdot a_i) \geqslant 0.$$

By iteratively swapping opportunities, we can transform $\vec{S}$ into $\vec{S}^{\text{SB}}$, and by iteratively applying this inequality every time we swap opportunities, we can show that the summation in Equation (56) under $\vec{S}$ cannot exceed the value of that same summation under $\vec{S}^{\text{SB}}$. This is a contradiction. □

B.2. Score-based Rankings under Alternative Choice Models

Another popular choice model with position effects is the *cascade model*, introduced and analyzed in Kempe and Mahdian (2008) and Aggarwal et al. (2008). In this model, the probability of selecting an item depends on all of the better-ranked items. Nevertheless, Kempe and Mahdian (2008) and Aggarwal et al. (2008) show that for a fixed set of ranked items, a score-based ranking (akin to $\vec{S}^{\text{SB}}$) maximizes the weighted sum of expected clicks.

We begin by defining a cascade choice model in our context for a representative volunteer t.

[15] Ranking in descending order obtains good performance under our assumption that top ranks are more appealing, i.e., θ_k is non-increasing in the position k. If this assumption is violated, we can follow the more general strategy of mapping the opportunity with the k^{th} highest score into the k^{th} most appealing position.

Definition 5 (Cascade Model) *For a given volunteer t, the cascade model depends on two sets of features: (i) a quality score for each opportunity j, denoted $q_{j,t}$, and (ii) a continuation probability for each opportunity j, denoted $c_{j,t}$, where we assume $c_{j,t}+q_{j,t}\leqslant 1$ to ensure that each volunteer chooses at most one opportunity. The volunteer's choice probability for opportunity i under ranking $\vec{S}$ is given by*

$$\phi_{i,t}(\vec{S}) \quad = \quad q_{i,t}\prod_{k=1}^{k(\vec{S},i)-1} c_{k^{-1}(\vec{S},k),t}, \tag{57}$$

where $k(\vec{S},i)$ denotes the position of opportunity i and $k^{-1}(\vec{S},k)$ denotes the opportunity in position k under ranking $\vec{S}$.

In words, a volunteer t searching under the cascade model will first consider the top ranked opportunity. For each opportunity i that they consider, they will do exactly one of three things: (i) sign up, which happens with probability $q_{i,t}$, (ii) continue to the next opportunity, which happens with probability $c_{i,t}$, or (iii) exit without signing up, which happens with probability $1-q_{i,t}-c_{i,t}$. This corresponds to the marginal probability distribution given by Equation (57).[16]

Given a set of opportunities to rank, Kempe and Mahdian (2008) and Aggarwal et al. (2008) show that the optimal solution to (PRO) under the cascade model is to rank the opportunities in descending order of $\psi(Z_{i,t-1}/m_i)\cdot q_{i,t}/(1-c_{i,t})$. The first term is the penalty function, while we can view the ratio $q_{i,t}/(1-c_{i,t})$ as the appeal of each opportunity. Specifically, if we define $a_{i,t}:=q_{i,t}/(1-c_{i,t})$, then the score-based ranking $\vec{S}^{\text{SB}}$ defined in Proposition 2 is an optimal solution to the optimization problem (PRO).

Going beyond the specific classes of MNL and cascade models, a score-based ranking like $\vec{S}^{\text{SB}}$ is exactly optimal under the general class of *click-through-rate (CTR) models* considered in L'Ecuyer et al. (2017), where each item's CTR depends on its position, its relevance, and the relevance of other items. Crucially, in this class of choice models, the probability of selecting an item *does not* depend on the position of other items. Intuitively, if placing an item in a "good" position does not negatively impact other items (aside from the fact that the position is now filled), then greedy score-based ranking strategies perform well. The strong performance of greedy rankings is also discussed in Ferreira et al. (2022).

B.3. Performance of Score-Based Ranking under Misspecification

In this section, we establish the following guarantee for the performance of score-based rankings when opportunities' appeals are mis-estimated.

Proposition 3 (Performance of Score-Based Ranking under Misspecification) *Suppose volunteer t follows an MNL choice model with position bias (see Definition 4). Furthermore, suppose $\bar{p}_t$ is an upper bound on volunteer t's* aggregate *sign-up probability under any ranking, i.e., $\bar{p}_t\geqslant \max_{\vec{S}\in\mathcal{S}}\sum_{i\in\vec{S}}\phi_{i,t}(\vec{S})$ and suppose there is an estimated vector of appeals $\hat{\mathbf{a}}$ and a true vector of appeals $\mathbf{a}$, where $a_i\in[\nu\cdot\hat{a}_i,\hat{a}_i/\nu]$ for all*

[16] The original cascade models do not specify a correlation between the "sign-up" and "continue" decisions. We assume a perfect negative correlation, to fit our model where volunteers sign up for at most one opportunity. This change does not impact the cascade model's selection probabilities or analysis; see Remark 2.1 in Kempe and Mahdian (2008) or Footnote 3 in Aggarwal et al. (2008) for further discussion.

$i \in [n]$ and some $\nu \in (0,1]$. Then the score-based ranking $\vec{S}^{\text{SBM}}$ that ranks opportunities in descending order of $\psi(Z_{i,t-1}/m_i) \cdot \hat{a}_{i,t}$ satisfies

$$\sum_{i \in \vec{S}^{\text{SBM}}} \psi(Z_{i,t-1}/m_i) \cdot \phi_{i,t}(\vec{S}^{\text{SBM}}) \quad \geqslant \quad (1-\bar{p}_t) \cdot \nu^2 \cdot \max_{\vec{S} \in \mathcal{S}} \sum_{i \in \vec{S}} \psi(Z_{i,t-1}/m_i) \cdot \phi_{i,t}(\vec{S}) \tag{58}$$

In words, $\vec{S}^{\text{SBM}}$ provides a $(1-\bar{p}_t) \cdot (\nu^2)$-approximation for the optimization problem (PRO).

When combined with Corollary 2, Proposition 3 establishes a guarantee on the performance of a score-based ranking under the MNL model with position bias even in the presence of misspecification. Comparing to Proposition 2, we see that appeal misspecification by a factor of ν impacts the approximation guarantee by a factor of ν^2.

Proof of Proposition 3 The proof follows nearly identical steps to the proof of Proposition 2. The only difference comes in the inequality in Line (54) in Step 2 of the proof, which no longer holds due to the presence of misspecification. Instead, we replace Step 2 of that proof with the following series of inequalities, which hold for any ranking $\vec{S}$.

$$\begin{aligned}
\sum_{i \in \vec{S}} \psi_i \cdot \phi_i(\vec{S}) \quad &= \quad \frac{\sum_{i \in \vec{S}} \psi_i \cdot \theta_{k(\vec{S},i)} \cdot a_i}{1 + \sum_{i \in \vec{S}} \cdot \theta_{k(\vec{S},i)} \cdot a_i} \\
&\leqslant \quad \frac{\sum_{i \in \vec{S}} \psi_i \cdot \theta_{k(\vec{S},i)} \cdot a_i}{1} \\
&\leqslant \quad \frac{\sum_{i \in \vec{S}} \psi_i \cdot \theta_{k(\vec{S},i)} \cdot \hat{a}_i}{\nu} && (59) \\
&\leqslant \quad \frac{\sum_{i \in \vec{S}^{\text{SBM}}} \psi_i \cdot \theta_{k(\vec{S}^{\text{SBM}},i)} \cdot \hat{a}_i}{\nu} && (60) \\
&\leqslant \quad \frac{\sum_{i \in \vec{S}^{\text{SBM}}} \psi_i \cdot \theta_{k(\vec{S}^{\text{SBM}},i)} \cdot a_i}{\nu^2} && (61) \\
&\leqslant \quad \frac{\sum_{i \in \vec{S}^{\text{SBM}}} \psi_i \cdot \theta_{k(\vec{S}^{\text{SBM}},i)} \cdot a_i}{M(1-\bar{p})\nu^2}
\end{aligned}$$

Inequality in (59) comes from the assumption that $a_i \leqslant \hat{a}_i/\nu$ for each opportunity i. We then apply Claim 2 to establish Line (60), which holds because $\vec{S}^{\text{SBM}}$ is a score-based ranking using appeals $\hat{\mathbf{a}}$. We then return to the true appeals in Line (61) using the inequality $a_i \geqslant \nu \cdot \hat{a}_i$, which comes from the statement of the proposition.

From here, we can follow the same path as the proof of Proposition 2 to show that for any ranking $\vec{S}$,

$$\sum_{i \in \vec{S}^{\text{SBM}}} \psi_i \cdot \phi_i(\vec{S}^{\text{SBM}}) \quad \geqslant \quad \frac{\sum_{i \in \vec{S}^{\text{SBM}}} \psi_i \cdot \theta_{k(\vec{S}^{\text{SBM}},i)} \cdot a_i}{M} \quad \geqslant \quad (1-\bar{p}) \cdot \nu^2 \cdot \sum_{i \in \vec{S}} \psi_i \cdot \phi_i(\vec{S})$$

This completes the proof of Proposition 3. □

B.4. Discussion on the Impact of SmartSort's Penalty Parameter

In this section, we first establish the monotonicity of efficiency and equity with respect to c_1 for a single period under the MNL choice model with position bias. Then, we show that monotonicity may not hold when considering a longer time horizon due to the interaction between the sign-up history and the current ranking.

Proposition 4 (Monotonicity with Respect to c_1) *Consider a single volunteer t and a vector $\mathbf{Z}_{t-1} \in \mathbb{Z}_+^n$ consisting of the sign-up history for each opportunity. If volunteer t follows an MNL choice model with position bias, then the efficiency (resp. equity) under SmartSort after period t is non-increasing (resp. non-decreasing) in c_1, fixing $c_2 = 1$.*

Proof of Proposition 4 To prove Proposition 4, it is sufficient to show that the probability of *any* sign-up in period t is non-increasing in c_1, and the probability of a sign-up for an opportunity without a prior sign-up is non-decreasing in c_1. Combined with the fixed sign-up history, this would imply that efficiency (resp. equity) is non-increasing (resp. non-decreasing) in c_1.

The proof relies on the following key idea: increasing the penalty c_1 can never move a more appealing opportunity above a less appealing opportunity in the ranking. To formalize this idea, we will use $\vec{S}^{\mathrm{SS}(c_1)}$ to denote the ranking under SmartSort with parameter c_1. Furthermore, $k(\vec{S}^{\mathrm{SS}(c_1)}, i)$ denotes the position of opportunity i under SmartSort with parameter c_1.

Claim 3 (Impact of c_1 on Relative Ranking of Two Opportunities) *Consider two opportunities i and j and two parameters c_1' and c_1'', with $c_1' < c_1''$. Suppose $\theta_{k(\vec{S}^{\mathrm{SS}(c_1')},i),t} > \theta_{k(\vec{S}^{\mathrm{SS}(c_1')},j),t}$, i.e., suppose i has a better ranking than j under SmartSort with parameter c_1'. If $\theta_{k(\vec{S}^{\mathrm{SS}(c_1'')},i),t} < \theta_{k(\vec{S}^{\mathrm{SS}(c_1'')},j),t}$, i.e., if j has a better ranking under SmartSort with a harsher penalty, then $a_{i,t} \geqslant a_{j,t}$.*

We can use Claim 3 (proved below) to establish monotonicity of efficiency and equity. The probability of obtaining any sign-up in period t under an MNL choice model with position bias (see Definition 4) is given by

$$\sum_{i \in \vec{S}^{\mathrm{SS}(c_1)}} \phi_{i,t}(\vec{S}^{\mathrm{SS}(c_1)}) \quad = \quad \frac{\sum_{i \in \vec{S}^{\mathrm{SS}(c_1)}} \theta_{k(\vec{S}^{\mathrm{SS}(c_1)},i),t} \cdot a_{i,t}}{1 + \sum_{j \in \vec{S}^{\mathrm{SS}(c_1)}} \theta_{k(\vec{S}^{\mathrm{SS}(c_1)},j),t} \cdot a_{j,t}}, \tag{62}$$

which is increasing in $\sum_{i \in \vec{S}^{\mathrm{SS}(c_1)}} \theta_{k(\vec{S}^{\mathrm{SS}(c_1)},i),t} \cdot a_{i,t}$. This summation is maximized by ranking opportunities in descending order of appeal (see Claim 2).

By Claim 3, any difference in the ranking that results from increasing c_1 can be obtained by iteratively swapping opportunities, where the higher-appeal opportunity moves to a worse ranking. Each of these shifts weakly reduces $\sum_{i \in \vec{S}^{\mathrm{SS}(c_1)}} \theta_{k(\vec{S}^{\mathrm{SS}(c_1)},i),t} \cdot a_{i,t}$, and hence cannot improve the probability of obtaining a sign-up. To see this more formally, suppose that under ranking $\vec{S}$, opportunity i has a better rank than opportunity j and opportunity i has higher appeal, i.e., $\theta_{k(\vec{S},i),t} \geqslant \theta_{k(\vec{S},j),t}$ and $a_{i,t} \geqslant a_{j,t}$. If we swap the ranks of opportunities i and j, the summation in Equation (56) will change by

$$\theta_{k(\vec{S},i),t} \cdot a_{j,t} - \theta_{k(\vec{S},j),t} \cdot a_{j,t} + \theta_{k(\vec{S},j),t} \cdot a_{i,t} - \theta_{k(\vec{S},i)} \cdot a_{i,t} \quad = \quad (\theta_{k(\vec{S},i),t} - \theta_{k(\vec{S},j),t})(a_{j,t} - a_{i,t}) \quad \leqslant \quad 0.$$

This proves that efficiency is non-increasing in c_1.

To establish that equity is non-decreasing in c_1, let us define the set $\mathcal{A}$ as the set of opportunities without a sign-up. Under an MNL choice model, the probability of any of those opportunities receiving a sign-up is given by:

$$\sum_{i \in \mathcal{A}} \phi_{i,t}(\vec{S}^{\mathrm{SS}(c_1)}) \quad = \quad \frac{\sum_{i \in \mathcal{A}} \theta_{k(\vec{S}^{\mathrm{SS}(c_1)},i),t} \cdot a_{i,t}}{1 + \sum_{j \in \vec{S}^{\mathrm{SS}(c_1)}} \theta_{k(\vec{S}^{\mathrm{SS}(c_1)},j),t} \cdot a_{j,t}} \tag{63}$$

Note that an increase in c_1 weakly reduces opportunities' scores under SmartSort but does not impact the scores of opportunities in $\mathcal{A}$. Thus, the rank of opportunities in $\mathcal{A}$ cannot get worse, and hence the numerator in the right hand side of (63) is non-decreasing in c_1. Furthermore, as previously shown, the denominator is non-increasing in c_1. Thus, the probability of an opportunity in $\mathcal{A}$ receiving a sign-up (equivalently, equity) is non-decreasing in c_1.

This completes the proof of Proposition 4. □

Proof of Claim 3 Consider two opportunities i and j, with $a_{i,t} > a_{j,t}$. We will show that an increase in c_1 can never cause opportunity i to move ahead of opportunity j in SmartSort's ranking.

Case (i): $Z_{i,t-1} \leqslant Z_{j,t-1}$. If opportunity i has received no more sign-ups than opportunity j, then i will have a better ranking under SmartSort than j for any c_1 based on Equation (8), as the score is weakly increasing in appeal and weakly decreasing in the number of sign-ups.

Case (ii): $Z_{i,t-1} > Z_{j,t-1}$. For this case, we limit our attention to $c_1 \leqslant 1/Z_{i,t-1}$, otherwise the score for opportunity i is equal to 0 (the minimum possible score) and it will be ranked below opportunity j.[17] Consider the following function of c_1 which is proportional to the difference in their scores under SmartSort:

$$g(c_1) \;=\; a_{i,t} \cdot (1 - e^{-1+c_1 \cdot Z_{i,t-1}}) - a_{j,t} \cdot (1 - e^{-1+c_1 \cdot Z_{j,t-1}}) \tag{64}$$

We will show that $g(c_1)$ can cross 0 at most once, which means that the opportunities can switch their relative rank at most once. Since $g(0) > 0$, i.e., i has a better rank when $c_1 = 0$, the only possible switch involves opportunity j moving to a better rank than opportunity i.

To see that $g(c_1)$ can cross 0 at most once, note that if $g(c_1) \leqslant 0$, we have $a_{i,t} \cdot (1 - e^{-1+c_1 \cdot Z_{i,t-1}}) \leqslant a_{j,t} \cdot (1 - e^{-1+c_1 \cdot Z_{j,t-1}})$. Rearranging terms and using the fact the $a_{i,t} > a_{j,t}$, we get

$$a_{j,t} \cdot e^{-1+c_1 \cdot Z_{j,t-1}} \;\leqslant\; a_{j,t} - a_{i,t} + a_{i,t} \cdot e^{-1+c_1 \cdot Z_{i,t-1}} \;\leqslant\; a_{i,t} \cdot e^{-1+c_1 \cdot Z_{i,t-1}} \tag{65}$$

Thus, if $g(c_1) \leqslant 0$, then

$$\begin{aligned} g'(c_1) \;&=\; -a_{i,t} \cdot Z_{i,t-1} \cdot e^{-1+c_1 \cdot Z_{i,t-1}} + a_{j,t} \cdot Z_{j,t-1} \cdot e^{-1+c_1 \cdot Z_{j,t-1}} \\ &\leqslant\; -a_{i,t} \cdot Z_{i,t-1} \cdot e^{-1+c_1 \cdot Z_{i,t-1}} + a_{i,t} \cdot Z_{j,t-1} \cdot e^{-1+c_1 \cdot Z_{i,t-1}} \qquad (66) \\ &=\; a_{i,t} \cdot e^{-1+c_1 \cdot Z_{i,t-1}} (-Z_{i,t-1} + Z_{j,t-1}) \\ &<\; 0 \end{aligned}$$

Inequality (66) comes from Line (65), and the final inequality comes from the fact that we are in case (ii), where $Z_{i,t-1} > Z_{j,t-1}$. This result implies that once $g(c_1) = 0$, it is monotonically decreasing (and hence will never cross 0 again). This completes the proof of Claim 3. □

Proposition 4 establishes that efficiency and equity are monotonic in c_1 within a single period. One might expect this relationship to hold when aggregating over multiple periods, which is consistent with what we observe in our simulations; however, there can be exceptions due to the impact of a sign-up on future rankings.

[17] If both opportunities are tied, we follow the convention of breaking ties in favor of opportunity j (the opportunity with lower appeal).

Example 2 *Consider an instance with three volunteers and two opportunities and suppose the first volunteer has already signed up for opportunity* 1*. The second and third volunteers follow MNL choice models with position bias where* $\theta_{1,t} = 1$ *and* $\theta_{2,t} = 0$*, i.e., they will only consider the top-ranked opportunity* i *and they will sign-up with probability* $a_{i,t}/(1 + a_{i,t})$*. The appeal vector for the second volunteer is* $\mathbf{a_2} = (2, 1)$ *and the appeal vector for the third volunteer is* $\mathbf{a_3} = (2, 0.5)$*. We will compare efficiency when* $c'_1 = 0.5$ *and* $c''_1 = 0.7$*.*

Under $c'_1 = 0.5$*, SmartSort ranks opportunity* 1 *in the top position for volunteer* 2 *and receives a sign-up with probability* 2/3*. If it does not receive a sign-up, then it will again rank opportunity* 1 *in the top position for volunteer* 3 *and receives a sign-up with probability* 2/3*. Otherwise, it will rank opportunity* 2 *in the top position and will receive a sign-up with probability* 1/3*. Overall, expected efficiency is* $1 + 2/3 + (2/3) \cdot (1/3) + (1/3) \cdot (2/3) = 2.11$*.*

Under $c''_1 = 0.7$*, SmartSort ranks opportunity* 2 *in the top position for volunteer* 2 *and receives a sign-up with probability* 1/2*. Regardless of the outcome, it will rank opportunity* 1 *in the top position for volunteer* 3 *and will obtain a sign-up with probability* 2/3*. Overall, expected efficiency is* $1 + 1/2 + 2/3 = 2.17$*.*

This example shows that a larger value of c_1 *can* lead to improved efficiency in aggregate.

B.5. Additional Details about Instances and Benchmarks in Figure 4

Here we describe the settings for the two example instances presented in Figure 4.

Market Size: Both instances consist of 200 opportunities and 500 homogeneous volunteers.

Opportunity Appeals: Since volunteers are homogeneous, each opportunity i has a universal appeal of a_i. We assume that a_i is drawn independently from a beta distribution which takes two parameters α and β.

- Instance 1: $\alpha = 1$, and $\beta = 0.75$
- Instance 2: $\alpha = 0.25$, and $\beta = 1$

In the left two panels of Figure 4, we present the distribution of opportunity appeals corresponding to each instance. We note that in the first example (shown in green), there are many high-appeal opportunities, whereas in the second example (shown in red), there are relatively few high-appeal opportunities.

Volunteer Choice: The probability that a volunteer signs-up for an opportunity is given by an MNL choice model with position bias (see Definition 4) with $\theta_k = (1 - k/11)^+/49.5$. Specifically, suppose opportunity i is ranked in position $k(\vec{S}, i)$. Then, we assume that the probability that volunteer t selects opportunity i when shown display ranking $\vec{S}$ is given by

$$\phi_{i,t}(\vec{S}) \quad = \quad \frac{\Big(1 - k(\vec{S}, i)/11\Big)^+ \cdot a_i}{49.5 + \sum_{j=1}^{200} \Big(1 - k(\vec{S}, j)/11\Big)^+ \cdot a_j}. \tag{67}$$

For this volunteer choice model, only opportunities ranked in the top ten positions have positive probability of a sign-up (capturing, e.g., the impact of being on the first page of search results), and the total probability of receiving any sign-up is at most 1/10 (which could be achieved if all opportunities had the maximum appeal of 1). By Proposition 2, a score-based ranking like the one used by SmartSort obtains at least a 0.9-approximation of (PRO).

Number of Simulation Iterations: For each instance and each considered value of c_1, i.e., for $c_1 \in \{0, 0.1, 0.2, \ldots, 1.0\}$, we conduct 100 iterations and report the average efficiency and equity across those simulations.

Benchmarks: We compare the performance of SmartSort against two different benchmarks.

- AO, the algorithm with the asymptotically-optimal competitive ratio as defined in Corollary 1. To generate the frontier of achievable combinations of equity and efficiency under this benchmark, we simulate outcomes for 11 different efficiency weights, $\gamma \in \{0, 0.1, \ldots, 1\}$. To solve (PRO), we leverage results in Abeliuk et al. (2016) and Rusmevichientong et al. (2020) to design a polynomial-time solution.
- RR, a simple round-robin algorithm which is a convex combination of an efficiency-oriented ranking and an equity-oriented ranking. With probability q, RR presents a greedy appeal-ordered ranking, and with probability $1-q$ it presents an equity-oriented ranking with the goal of providing equal exposure. Specifically, the equity-oriented ranking selects a permutation of opportunities at the outset of each simulation, and it rotates the opportunities after presenting a ranking by shifting the permutation 11 positions. (The choice of 11 ensures that volunteers consider a new set of 10 opportunities after each rotation and also ensures that each of the 200 opportunities would appear in each position once before appearing in a position for the second time.) To generate the frontier of achievable combinations of equity and efficiency under this benchmark, we simulate outcomes for 11 different convex combinations of the two rankings, $q \in \{0, 0.1, \ldots, 1\}$.

B.6. Simulations in Other Instances

Here we present similar results as in Figure 4, but in a range of different instances. First, in Figure 5, we present eight different variations on our base instances. In each variation, the relative distribution of opportunity appeals remains the same (shown in the left panel of Figure 4). We use the same three algorithms and parameters; for each one, we conduct 100 simulations and plot the average efficiency and equity that form the achievable frontier.

At a high level, we note that the key insights from our base simulations have not changed: (i) SmartSort continues to perform well in comparison to the two benchmarks, (ii) increasing c_1 allows us to trade off between equity and efficiency, and (iii) the trade-off is steeper in Instance 2, where there are fewer high-appeal opportunities. We now describe each variation in detail.

In the top row of Figure 5, we adjust the number of volunteers and opportunities. On the left, we have increased the number of volunteers from 500 to 1000. On the right, we have reduced the number of opportunities from 200 to 100. In both cases, the ratio of high-appeal opportunities to volunteers has decreased; as a consequence, we observe that the trade-off between equity and efficiency becomes slightly steeper in comparison to Figure 4.

In the second row of Figure 5, we adjust the parameters of the MNL choice model. On the left, we consider a version where the no-choice probability is limited. Specifically, we replace 49.5 in Equation (67) with 5.5, such that the maximum probability of receiving a sign-up is now 0.5 instead of 0.1. (We have correspondingly adjusted the number of volunteers downward by a factor of 5 so that the maximum efficiency remains

similar.) Somewhat surprisingly, the gap between SS and AO shrinks, despite the fact that score-based rankings provide a worse guarantee when the maximum sign-up probability is larger (see Proposition 2). On the right, we consider an extreme version of position bias where the volunteer only considers the top-ranked opportunity i and signs-up with probability $a_i/(10+a_i)$. Despite the significant changes to the choice model parameters, the only significant change to the achievable frontiers is an increase in the range of attainable equity values.

In the third row of Figure 5, we consider varying levels of misspecification in the appeal distribution. In particular, we assume that each of the three algorithms misestimates the appeal of opportunity i by a multiplicative factor ν_i. On the left, we assume $\nu_i \sim \text{Unif}(0.75, 1.25)$. On the right, we assume $\nu_i \sim \text{Unif}(0, 2)$. Despite the possibility of significant misspecification, we observe minimal loss in performance.

Finally, in the bottom row of Figure 5, we consider a cascade choice model (see Definition 5) instead of an MNL choice model with position bias. Under a cascade model an opportunity's appeal is given by the ratio of its sign-up probability q_i and 1 minus its continuation probability c_i. We maintain the appeal distribution for each instance, which leaves us with one degree of freedom in choosing the q_i and c_i parameters. In the left (resp. right) panel, we set $c_i = 0.5$ (resp. $c_i = 0.9$) for all opportunities. The larger value for c_i in the right panel means that volunteers are more likely to browse further into the ranking, which leads to a larger value of equity even when focused purely on efficiency. Notably, the absence of a dotted red curve in the bottom right panel indicates that equity cannot be improved by mixing the greedy appeal-ordered ranking with a round-robin algorithm. In other words, efficiency and equity are *both* maximized under the RR algorithm when all of the weight is on the greedy appeal-ordered ranking.

Moving beyond the homogeneous volunteer setting used to generate Figures 4 and 5, in Figures 6 and 7 we present simulations where the appeal of an opportunity varies based on volunteer characteristics.

In Figure 6, we construct two instances where volunteers and opportunities have a single-dimensional type (e.g., a location). The appeal of opportunity i with type θ_i to volunteer t with type θ_t is given by $a_{i,t} = \left(1 - |\theta_t - \theta_i|\right)^+$. For each opportunity i, we draw a type θ_i from a Normal distribution with mean 0 and standard deviation 0.2 (in both instances). In Instance 1, for each volunteer t, $\theta_t \sim \mathcal{N}(0.4, 0.2)$, whereas in Instance 2, $\theta_t \sim \mathcal{N}(0.8, 0.2)$. In the left two panels of Figure 6, we display the *average* appeal of each opportunity, where the average is taken over all volunteers. In Instance 1, the opportunities and volunteers are more likely to overlap in the type space: some opportunities are appealing to some volunteers, while other opportunities are appealing to other volunteers. Hence, many different opportunities are reasonably appealing on average. By contrast, in Instance 2 the same small set of opportunities – those with higher types – are consistently the most appealing. Aside from the difference in appeals, the instances are identical to the instances shown in Figure 4 and described in Appendix B.5.

In the right panel of Figure 6, we plot equity and efficiency for the same three algorithms and parameters, averaging over 100 simulations in each case. The results are consistent with our prior simulations involving homogeneous volunteers: (i) SmartSort performs well in comparison to the two benchmarks, (ii) increasing c_1 allows us to trade off between equity and efficiency, and (iii) the trade-off is slightly steeper in Instance 2, where there are fewer opportunities with high average appeal. One substantial difference between Instance

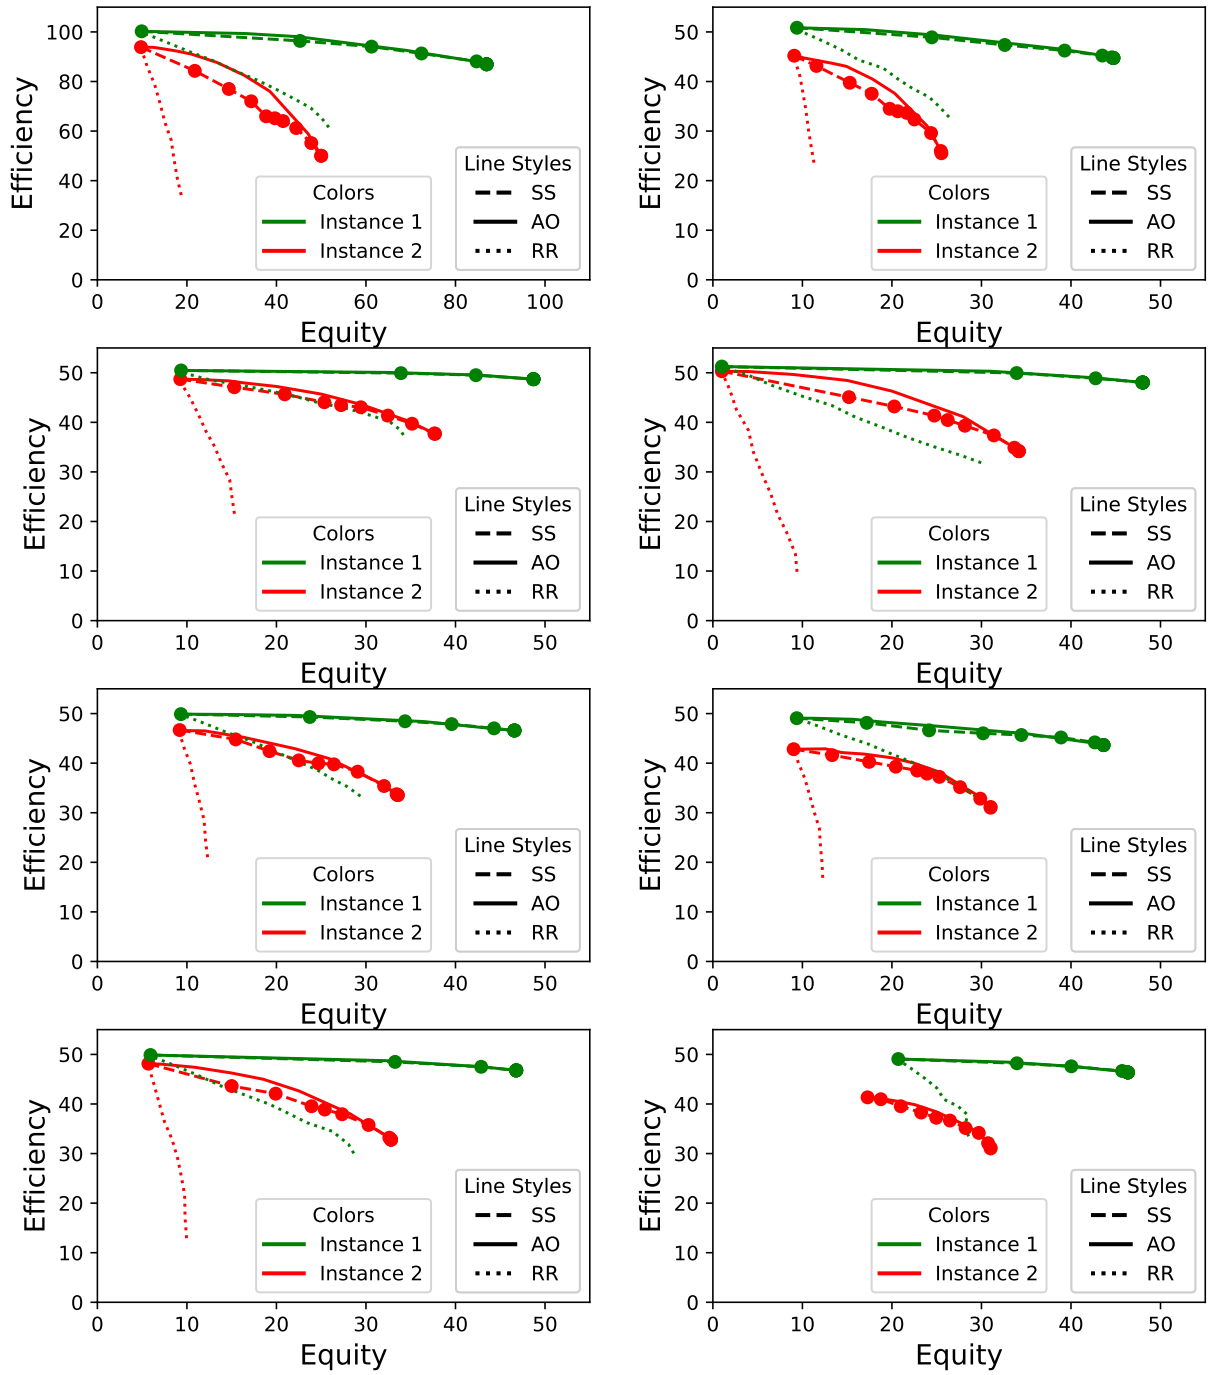

Figure 5 Efficiency and equity in variations of our base instances. The dashed curves show the efficiency and equity outcomes under SS for both instances, varying the value for the penalty parameter c_1. The results shown are averaged across 100 simulations; the standard errors are illustrated by the size of the ovals. The solid (resp. dotted) curves indicate the frontier of efficiency and equity achievable by AO (resp. RR) for various values of the efficiency weight γ (resp. for various probabilities of presenting a greedy ranking).

1 and our prior simulations is the fairly high degree of equity even when focusing entirely on efficiency. This

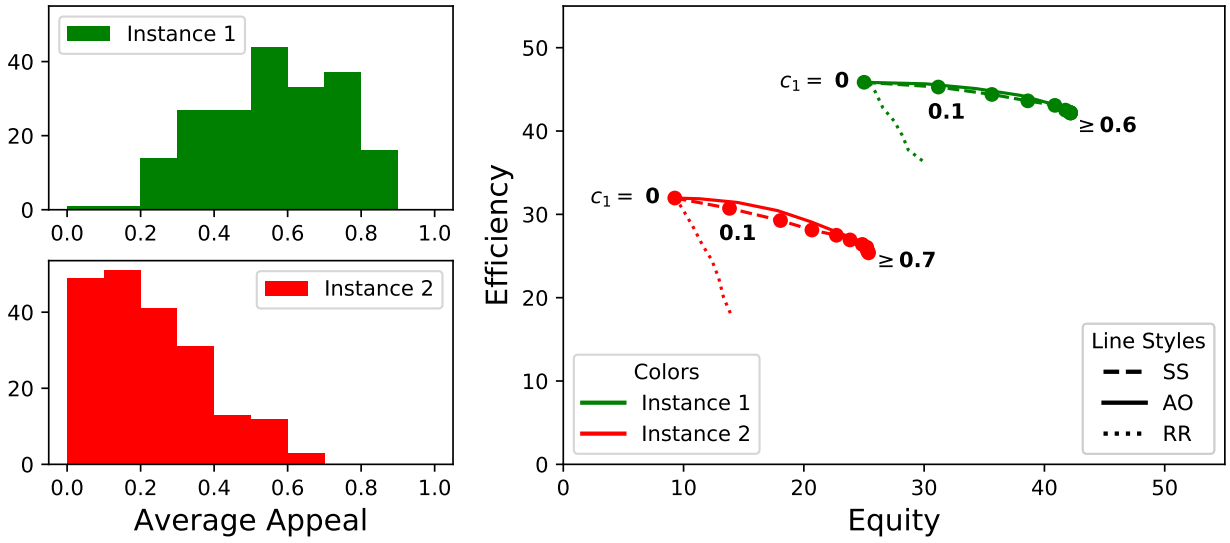

Figure 6 The left two panels show the distributions of opportunities' average appeals in the two instances where opportunities and volunteers have one-dimensional types. The dashed curves in the right panel show the efficiency and equity outcomes under SS for both instances, varying the value for the penalty parameter c_1. The results shown are averaged across 100 simulations; the standard errors are illustrated by the size of the ovals. The solid (resp. dotted) curves indicate the frontier of efficiency and equity achievable by AO (resp. RR) for various values of the efficiency weight γ (resp. for various probabilities of presenting a greedy ranking).

stems from the fact that volunteers preferences are heterogeneous: we get some additional equity just by catering to their diverse preferences.

In Figure 7, we construct two instances based on the DFW marketplace on VM, each with 300 opportunities and 3000 volunteers. In Instance 1, we determine the appeal between an opportunity i and a volunteer t using SmartSort's formula (see Equation 9), which requires us to specify (i) a location for opportunity i, (ii) a recency for opportunity i, and (iii) a location for volunteer t. For each opportunity, we generate a location by sampling the zip code of a randomly-chosen opportunity on VM in the DFW area. The opportunity's recency is chosen uniformly at random between 0 and 100. For each volunteer, we generate a location by sampling a zip code proportional to the population distribution in the DFW area. If an opportunity meets the criteria for appearing in the display ranking (i.e., if it is within 20 miles of the volunteer and has a recency of at most 90 days), we use Equation (9) to determine its appeal; otherwise, we assume its appeal is 0. We continue to use the same MNL choice model with position bias, given in Equation (67).

Instance 2 is identical to Instance 1, except that we allow for potential inaccuracy in the appeal estimates. In particular, we assume that the true appeal of opportunity i consistently differs from its estimated appeal under SmartSort by a factor ν_i, where $\nu_i \sim \text{Unif}[0.5, 1.5]$.

In the two left panels of Figure 7 we report the average appeal of each opportunity, where the average is taken over all volunteers. In Instance 1, a large number of opportunities have a similarly moderate average appeal. For this instance, opportunities in similar locations have nearly identical appeal toward each volunteer. In Instance 2, where the true appeal also includes an opportunity-specific multiplier, we observe a wider distribution of average appeals.

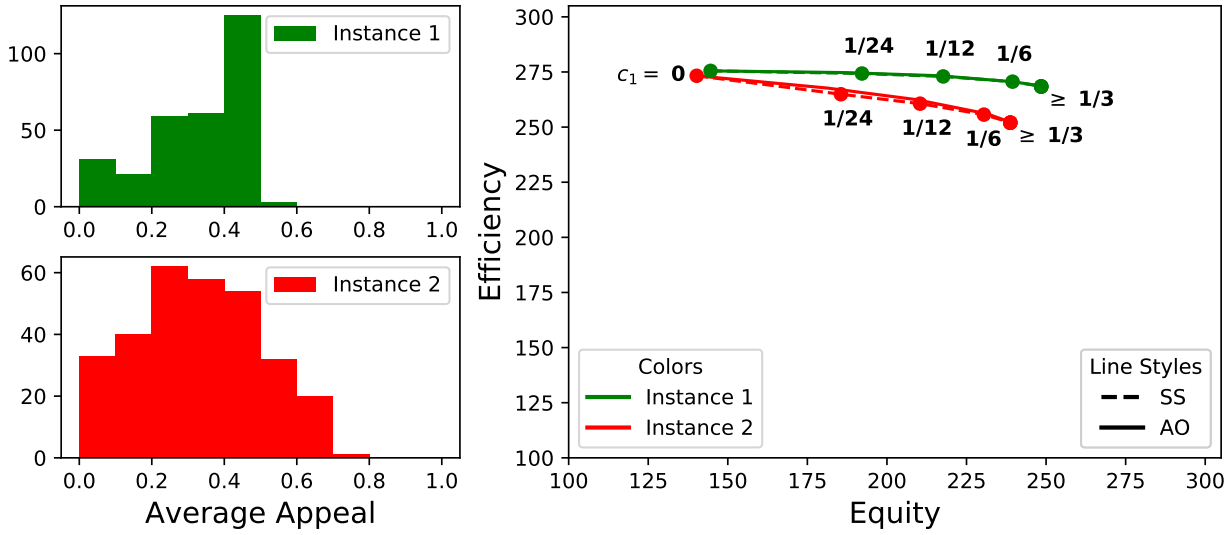

Figure 7 The left two panels show the distributions of opportunities' average appeals in the two instances where opportunities and volunteers have characteristics calibrated on data from DFW. The dashed curves in the right panel show the efficiency and equity outcomes under SS for both instances, varying the value for the penalty parameter c_1. The results shown are averaged across 100 simulations; the standard errors are illustrated by the size of the ovals. The solid curves indicate the frontier of efficiency and equity achievable by AO for various values of the efficiency weight γ.

In the right panel of Figure 7, we we plot the equity and efficiency attainable by SmartSort and our benchmark AO algorithm. (We do not plot the performance of our benchmark RR algorithm because equity cannot be improved by mixing with a round-robin algorithm.) To better assess the granular impacts of increasing c_1, we conduct simulations for $c_1 \in [0, 1/24, 1/12, 1/6, 1/3, 1]$. We observe that even small values of c_1 can lead to substantial gains in equity. In Instance 1, for $c_1 \leqslant 1/12$ we obtain these equity gains while reducing efficiency by less than 1%. In Instance 2, where appeals are misspecified, there is a greater risk to efficiency, though the drop is still less than 5% for $c_1 \leqslant 1/12$.

B.7. Additional Details on Calibration of c_1

To provide some intuition behind our rationale for the specific choice of 1/12, consider the following data analysis and thought experiment involving two zip codes in the Dallas-Forth Worth region. Zip code 75039 is at the outskirts of Dallas with few nearby opportunities. In contrast, zip code 75201 is at the center of Dallas where many opportunities are located. Suppose every day between March 15$^{\text{th}}$ and May 23$^{\text{rd}}$ a representative volunteer searches from zip code 75039 (resp. 75201).[18] The left (resp. right) panel of Figure 8 shows the average daily distribution of opportunity appeals — as defined in Equation (9) — for the 100 most appealing opportunities for this representative volunteer. These empirical distributions suggest considerable differences in the appeal distribution. To highlight this, in the same histograms, we also plot a vertical line to indicate the average score of the 25th most appealing opportunity.[19] We remind that each page on VM displays 25

[18] Note that Dallas-Forth Worth was the site of our first experiment and the time period March 15$^{\text{th}}$ to May 23$^{\text{rd}}$ coincides with its corresponding pre-experiment period. See Section 4.1 for more details.

[19] Note that the dashed lines show one standard deviation of this cutoff value and the red bars indicate the standard error of the height of each bar.

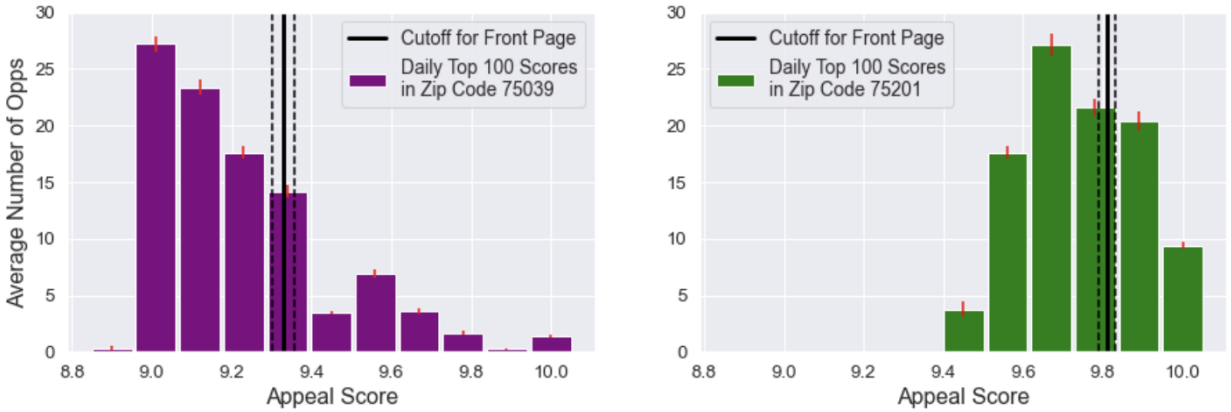

Figure 8 The left (resp. right) panel shows the average daily distribution of opportunity appeals — as defined in Equation (9) — for the 100 most appealing opportunities for a representative volunteer from zip code 75039 (resp. 75201). The solid black lines correspond to the average appeal score of the last opportunity to make the front page.

opportunities. Thus, in the absence of any penalty (i.e., if $c_1 = 0$), the black vertical line indicates the appeal "cutoff" for first page. This cutoff is considerably lower in zip code (75039) compared to 75201 (9.33 vs. 9.81).

Based on this data analysis, the following thought experiment contributed to our choice of $c_1 = 1/12$. Suppose initially, no opportunity has a sign-up. In zip code 75039 (resp. 75201) consider an opportunity o_i (resp. o_j) with the maximal appeal of 10. Without a sign-up, opportunity o_i (resp. o_j) would be ranked on top. (Let us ignore the possibility of ties for the ease of argument.) When $c_1 = 1/12$, one sign-up reduces the score of opportunity o_i (resp. o_j) from 10 to approximately 9.5, according to SmartSort's formula given by Equation (8). In zip code 75039, this would move opportunity o_i a few places down the display ranking, but it would likely remain on the first page. On the other hand, in zip code 75201, o_j would likely be moved off the front page; however due to the abundance of high-appeal opportunities, it will be replaced by an opportunity with a similar appeal of approximately 9.9 (i.e., an opportunity within a mile of the volunteer's search). Under the (conservative but reasonable) assumption that volunteers only pay attention to the first page of search results, we see that this level of penalization is small enough to keep the high-appeal opportunity on the first page in situations where there is a scarcity of high-appeal opportunities (i.e., zip code 75039) thus reducing the risk of efficiency loss. At the same time, it is sufficient to prioritize high-appeal opportunities with no sign-ups.

We now repeat the above thought experiment for a larger penalty, $c_1 = 1/5$. In zip code 75039 (resp. 75201) consider again an opportunity o_i (resp. o_j) with the maximal appeal of 10. Without a sign-up, opportunity o_i (resp. o_j) would be ranked on top. When $c_1 = 1/5$, one sign-up reduces the score of opportunity o_i (resp. o_j) to approximately 8.7, according to SmartSort's formula given by Equation 8. In zip code 75039, opportunity o_i would now likely move off the front page and would be replaced by an opportunity with a score of approximately 9.4, which could be an opportunity without a sign-up but 6 miles away (thus potentially much less appealing). This could have a substantial negative impact on efficiency, especially if volunteers preferences for close opportunities are non-linear. In contrast, this change would not have an impact in zip

code 75201, relative to the previous discussion: opportunity o_j would still move off the front page of search results and would still be replaced by an opportunity with a similar score of approximately 9.9 (i.e., an opportunity without a sign-up within a mile of the volunteer's search).

Appendix C: Supplemental Material for the Dallas Experiment

C.1. Balance Tables

As mentioned in Section 4.1, DFW and HOU are the first and second largest regions in Texas. This makes HOU a natural comparison city. Drawing on zip-level data from the US Census American Community Survey 2020, in Table 3 we show that in terms of population, demographic composition, and income level the regions included in our analysis are fairly comparable.[20] Additionally, using aggregate city-level data on volunteer activity on VM provided by Google Analytics we also confirm that in our pre-experiment period (defined in Section 4.1) the number of visitors and sessions (i.e., the periods of time during which visitors interact with the VM site) are on the same order. We note that this is the only available data on users and sessions, and it is not available on a more granular level (i.e., on a zip code level). Furthermore, it only comprises a subset of the full activity in the DFW and HOU regions.

Table 3 Balance Table - DFW and HOU

	DFW	HOU
Total Population	6,282,445	5,538,4552
Pct Population Female	50.41%	50.20%
Pct Population White	57.75%	51.54%
Median Age (years)	32.9	33.3
Median Income	$54,747	$53,600
Total Users (pre-experiment)	22,531	16,608
Total Sessions (pre-experiment)	28,421	21,814

C.2. Parallel Trends Assumption

The key identifying assumption for the DID analyses is that of parallel trends between the treatment and control groups. In this appendix we provide evidence in support of the parallel trend assumption between DFW and HOU. First, in Figure 9 we simply plot the trends for our two outcome metrics across the two regions. Recall that the pre-experiment period consists of 10 consecutive weeks from Tuesday March 15, 2022 (Week 11) - May 23, 2022 (Week 20) and the experiment period consists of 10 consecutive weeks starting from May 24, 2022 (Week 21) - August 1, 2022 (Week 30). For both metrics, we observe strong and reasonably consistent seasonal trends in the pre-experiment period, with a noticeable summertime boost in volunteer activity in both Texas cities.

[20] https://www.census.gov/newsroom/press-kits/2021/acs-5-year.html

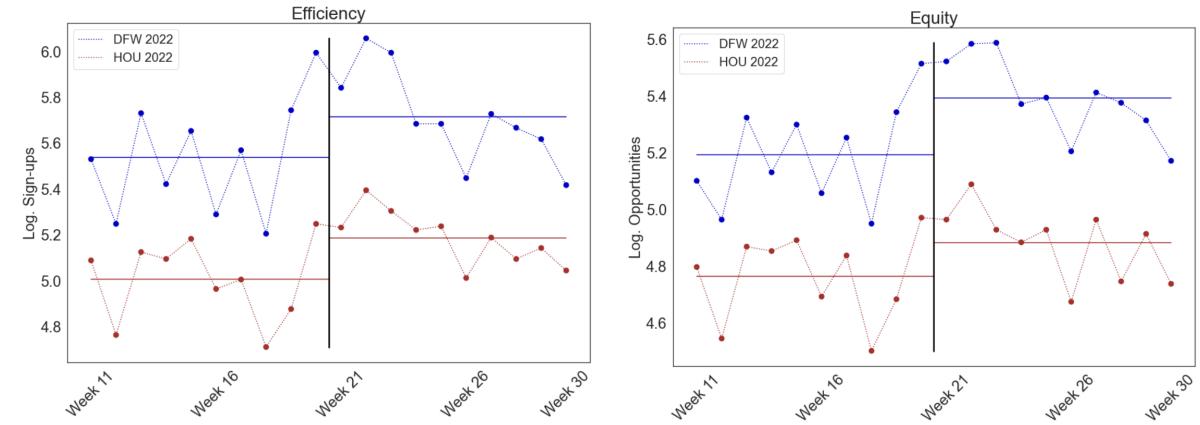

(a) (b)

Figure 9 The weekly seasonal trends of efficiency and equity in DFW and HOU.

To provide quantitative evidence for the parallel trend assumption, first we follow the approach of comparing the slope of time trends in the *pre-experiment* period across two regions (see Rios et al. 2022 for an example of this method). To that end, we consider the following specification:

$$\log(Y_{wc}) = \beta_0 + \beta_{\text{city}} \cdot \mathbf{1}_{\{c=\text{DFW}\}} + \delta_w + \rho_c w + \varepsilon_{wc}$$

In the above, $w \in [11, 20]$ is the week, $c \in \{\text{DFW}, \text{HOU}\}$ is the city, β_{city} is the city fixed effect, δ_w are week fixed effects, and ρ_c are the time slope in each of the control and treated regions.

We estimate the coefficient ρ_c for both DFW and HOU and use a t-test to determine if they are statistically significantly different from each other. The first two rows of Table 4 shows the estimated slopes for each of the outcome metrics in DFW and HOU. We see that there are significant seasonal effects pre-experiment, as both ρ_{DFW} and ρ_{HOU} are positive and statistically significant for both efficiency and equity. The final row shows the result of a t-test that considers the null hypothesis that $\rho_{\text{DFW}} = \rho_{\text{HOU}}$. The final row of Table 4 indicates that the difference between ρ_{DFW} and ρ_{HOU} is small for both efficiency and equity (0.024 and 0.022 respectively) and that both differences are statistically insignificant at the $\alpha = 0.05$ level. Thus, we cannot reject the null hypothesis that the slopes ρ_{DFW} and ρ_{HOU} are equal to each other for both efficiency and equity.

Finally, to provide further evidence for parallel trends, we conduct an event study (Pischke 2005). The purpose of this analysis is to consider whether there is a significant difference between the treated and the control group at any time pre-experiment. Concretely we consider the following specification:

$$\log(Y_{wc}) = \beta_{\text{city}} \cdot \mathbf{1}_{\{c=\text{DFW}\}} + \delta_w + \sum_{o=-9, o\neq 0}^{10} \beta_{w,\text{city}} \cdot \mathbf{1}_{\{w=20+o\}} \cdot \mathbf{1}_{\{c=DFW\}} + \varepsilon_{wc}, \tag{68}$$

In the above, $w \in [11, 30]$ is the week, $c \in \{\text{DFW}, \text{HOU}\}$ is the city, β_{city} is the city fixed effect, δ_w are week fixed effects, and $\beta_{w,\text{city}}$ are the leads and lags of the treatment. $\beta_{w,\text{city}}$ can be interpreted as the treatment effect if the treatment had occurred in that week, using the final pre-experiment week as a baseline (Pischke 2005).

Table 4 Comparison of Pre-trend Slopes in Log Outcomes in DFW vs HOU

	Efficiency	Equity
ρ_{DFW}	0.203***	0.183***
	(0.016)	(0.008)
ρ_{HOU}	0.178***	0.161***
	(0.010)	(0.005)
Difference in Slope	0.024	0.022
	(0.018)	(0.010)
T-Statistic	1.326	2.288
P-Value	0.221	0.051
Note:	*p<0.05; **p<0.01; ***p<0.001	

We plot the coefficients $\beta_{w,\text{city}}$ in Figure 10, where the bars represent a 95% confidence interval with bootstrapped standard errors calculated through resampling. We use the standard bootstrap method of sampling with replacements (see, e.g., Kuhn and Johnson 2013), which works in the following way to generate new instances of the treated and control cities:

1. We first define the total population of units as the opportunities receiving sign-ups in DFW over the 20 weeks of our analysis. Each of these units is an (opportunity, week, sign-ups) triplet corresponding to the number of volunteer sign-ups for that opportunity in that week. We then sample from this population with replacement to form a new set of (opportunity, week, sign-ups) triplets, such that this set is of the same length as the original population. We then do the same for the total population of opportunities receiving sign-ups in HOU over the 20 weeks of our analysis.[21]
2. We aggregate these resampled sets on the week level to generate a new instance of weekly outcomes in the treated and control cities. Then we estimate our coefficients $\beta_{w,\text{city}}$ according to Equation 68.
3. We repeat steps 1-2 for 1,000 iterations, creating a set of 1,000 estimates for $\beta_{w,\text{city}}$.
4. We calculate the standard error of $\beta_{w,\text{city}}$ over these iterations, and using these standard errors, we construct 95% confidence intervals for our estimated $\beta_{w,\text{city}}$ in Figure 10.

Our pre-trend analysis generally supports the parallel trends assumption for our data: we do not see a pattern of steady increase or decrease in the $\beta_{w,\text{city}}$ coefficients over the pre-experiment weeks, suggesting that the trends are neither diverging nor converging over time. In addition, across all metrics, only a few of the coefficients for the leads of the treatment are statistically distinguishable from 0, which suggests that the parallel trends assumption is reasonable.

C.3. Justification for Including Time-Varying Covariates in DID Analysis

Including a time-varying covariate, such as efficiency, to a difference-in-differences regression specification requires proper justification. Intuitively, if the weekly efficiency outcome explains some of the variance in the weekly equity outcome then not controlling for it could result in omitted variable bias, leading to imprecise

[21] We choose this population as the sampling population to avoid introducing bias into the outcome metrics. By construction, the total equity summed across all 20 weeks will be the same; furthermore, in expectation, total efficiency summed across all 20 weeks will also be the same.

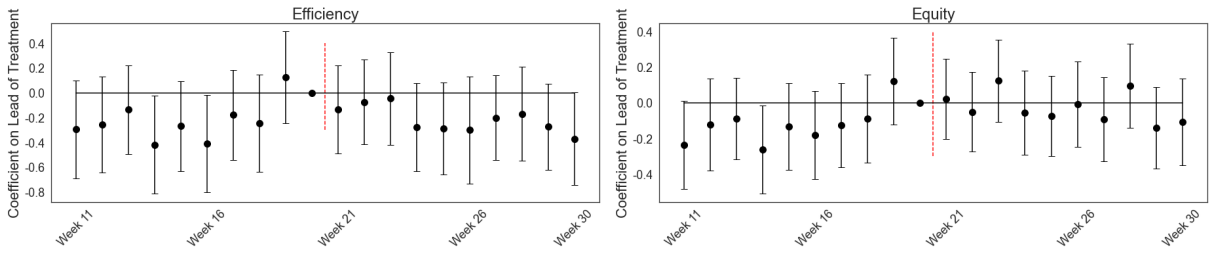

(a) (b)

Figure 10 An event study pre-trend analysis on the seasonal trends of our outcome metrics, with 95% confidence intervals based on bootstrapped standard errors.

measurement of the variance on the treatment coefficient. However, if efficiency is itself strongly correlated with the treatment, then including this variable would introduce co-linearity issues that could bias the interpretation of the treatment coefficient. Therefore, it would only be appropriate to control for the time-varying covariate of efficiency if efficiency and the equity outcome are highly correlated but efficiency and the treatment indicator are not correlated once we control for the other covariates. As mentioned in the main text, this is in fact the case in our data. We show the correlation values throughout this section.

To rigorously make the above point, let us first introduce new notations. First consider the base specification from Equation 10; let us refer to this as Specification A (SA).

$$\log(Y_{wc}) = \beta_0 + \beta_{\text{treat}} \cdot T_{wc} + \beta_{city} \cdot \mathbf{1}_{\{c=\text{DFW}\}} + \beta_{post} \cdot \mathbf{1}_{\{w\in[21,30]\}} + \varepsilon_{wc}, \tag{SA}$$

Here T_{wc} is the indicator variable for whether the treatment is turned on or not in week $w \in [11, 30]$ and city $c \in \{\text{DFW}, \text{HOU}\}$ and Y_{wc} is the equity outcome in week w in city c. Note that ε_{wc} is the residual error term and thus assumed to be uncorrelated with all covariates. By definition, we have:

$$\mathbb{E}[\varepsilon_{wc}^2] := 1 - R^2_{(\text{SA})}$$

Now by the Frisch-Waugh-Lovell Theorem (Frisch and Waugh (1933), Angrist and Pischke (2013)), we can write the treatment coefficient in (SA) as follows (note that this will be an unbiased estimator).

$$\beta_{treat} = \frac{Cov(Y_{wc}, \tilde{T}^A_{wc})}{Var(\tilde{T}^A_{wc})},$$

Here $\tilde{T}^A_{wc}$ is the residual error term of a linear regression that regresses T_{wc} on the city and post-treatment period indicators, $\mathbf{1}_{\{c=\text{DFW}\}}$ and $\mathbf{1}_{\{w\in[21,30]\}}$. This regression equation is shown below:

$$T_{wc} = \alpha_{city} \cdot \mathbf{1}_{\{c=\text{DFW}\}} + \alpha_{post} \cdot \mathbf{1}_{\{w\in[21,30]\}} + \tilde{T}^A_{wc} \tag{69}$$

Next consider the specification for $\log(Y_{wc})$ that additionally controls for the log-transformed weekly number of sign-ups as in Column 4 in Table 2. Let us refer to this as Specification B (SB), expressed as follows, which is equivalent to the specification introduced in Equation (12).

$$\log(Y_{wc}) = \gamma_0 + \gamma_{\text{treat}} \cdot T_{wc} + \gamma_{city} \cdot \mathbf{1}_{\{c=\text{DFW}\}} + \gamma_{post} \cdot \mathbf{1}_{\{w\in[21,30]\}} + \gamma_L \cdot L_{wc} + \mu_{wc}, \tag{SB}$$

In the above, L_{wc} represents the log-transformed number of sign-ups in week $w \in [11, 30]$ and city $c \in \{\text{DFW}, \text{HOU}\}$. Note that again μ_{wc} is the residual error and assumed to be uncorrelated with all covariates. By definition, we have:

$$\mathbb{E}[\mu_{wc}^2] := 1 - R^2_{(\text{SB})}$$

Applying the Frisch-Waugh-Lovell Theorem again, we can express the treatment coefficient in (SB) as follows [22]:

$$\gamma_{treat} = \frac{Cov(Y_{wc}, \tilde{T}^B_{wc})}{Var(\tilde{T}^B_{wc})},$$

where $\tilde{T}^B_{wc}$ is similarly the residual error term of a linear regression that regresses T_{wc} on L_{wc} as well as the city and post-treatment period indicators, $\mathbf{1}_{\{c=\text{DFW}\}}$ and $\mathbf{1}_{\{w\in[21,30]\}}$. This regression equation is shown below:

$$T_{wc} = \eta_{city} \cdot \mathbf{1}_{\{c=\text{DFW}\}} + \eta_{post} \cdot \mathbf{1}_{\{w\in[21,30]\}} + \eta_L \cdot L_{wc} + \tilde{T}^B_{wc} \tag{70}$$

Estimate of γ_{treat}:

We want to know how controlling for L_{wc} will impact the estimate of the treatment effect; specifically, we want to show that in our setting $\gamma_{treat} \approx \beta_{treat}$. To do this, it suffices to show that $Var(\tilde{T}^B_{wc}) \approx Var(\tilde{T}^A_{wc})$ and that $Cov(Y_{wc}, \tilde{T}^B_{wc}) \approx Cov(Y_{wc}, \tilde{T}^A_{wc})$. And to do this, it suffices to show that $\tilde{T}^B_{wc} \approx \tilde{T}^A_{wc}$.

Towards this goal, we first use the Frisch-Waugh-Lovell Theorem to express the coefficient on $L_{w,c}$ in (70):

$$\eta_L = \frac{Cov(T_{wc}, \tilde{L}_{wc})}{Var(\tilde{L}_{wc})}$$

Here $\tilde{L}_{wc}$ is the residual error term of a linear regression that regresses L_{wc} on the city and post-treatment period indicators, $\mathbf{1}_{\{c=\text{DFW}\}}$ and $\mathbf{1}_{\{w\in[21,30]\}}$. This regression equation is shown below:

$$L_{wc} = \zeta_{city} \cdot \mathbf{1}_{\{c=\text{DFW}\}} + \zeta_{post} \cdot \mathbf{1}_{\{w\in[21,30]\}} + \tilde{L}_{wc} \tag{71}$$

Notice that when the correlation between T_{wc} and $\tilde{L}_{wc}$ (i.e., $\rho_{T,\tilde{L}}$) is low and all covariates are finitely bounded, we can conclude $\eta_L \approx 0$. In our setting $\rho_{T,\tilde{L}} = -0.01$, so $\eta_L \approx 0$. Next we make the crucial observation that $\tilde{T}^A_{wc}$ and $\tilde{T}^B_{wc}$ are residual error terms and thus assumed to be uncorrelated with the covariates in their specifications. Thus, having $\eta_L \approx 0$ implies $\tilde{T}^B_{wc} \approx \tilde{T}^A_{wc}$. Having $\tilde{T}^B_{wc} \approx \tilde{T}^A_{wc}$ implies both that $Var(\tilde{T}^B_{wc}) \approx Var(\tilde{T}^A_{wc})$ and that $Cov(Y_{wc}, \tilde{T}^B_{wc}) \approx Cov(Y_{wc}, \tilde{T}^A_{wc})$.

Thus, $\rho_{T,\tilde{L}} \approx 0$ implies that $\beta_{treat} \approx \gamma_{treat}$; since both β_{treat} and γ_{treat} are unbiased estimators, we expect to observe that $\hat{\beta}_{treat} \approx \hat{\gamma}_{treat}$. Table 2 shows exactly this pattern — the coefficient in column 3 (0.082) is

[22] Again, by the same theorem, this is an unbiased estimator.

approximately equal to the coefficient in column 4 (0.086) — suggesting that controlling for L_{wc} is appropriate. The same argument holds for the specification in column 5. For this specification, $\rho_{T,\tilde{L}} = -0.013$ and the coefficient of 0.085 is approximately equal to that in column 3, suggesting that controlling for L_{wc} is appropriate here as well.[23]

In comparison, consider the case where the log-transformed number of sign-ups in a week is causally impacted by the treatment. In this case, the log-transformed number of sign-ups in a week will be strongly correlated with treatment, even after controlling for the city and post-treatment period indicators. That is, the correlation $\rho_{T,\tilde{L}} \not\approx 0$. In this case, $\tilde{T}^B_{wc} \not\approx \tilde{T}^A_{wc}$ and therefore $\beta_{treat} \not\approx \gamma_{treat}$. In this case, we can quantify how much the coefficients β_{treat} and γ_{treat} will differ by considering the size of the difference between $\tilde{T}^B_{wc}$ and $\tilde{T}^A_{wc}$. This difference will be given by:

$$\begin{aligned}\tilde{T}^B_{wc} - \tilde{T}^A_{wc} &= T_{wc} - \eta_{city} \cdot \mathbf{1}_{\{c=\text{DFW}\}} - \eta_{post} \cdot \mathbf{1}_{\{w\in[21,30]\}} - \eta_L \cdot L_{wc} - T_{wc} + \alpha_{city} \cdot \mathbf{1}_{\{c=\text{DFW}\}} + \alpha_{post} \cdot \mathbf{1}_{\{w\in[21,30]\}} \\ &= (\alpha_{city} - \eta_{city})\mathbf{1}_{\{c=\text{DFW}\}} + (\alpha_{post} - \eta_{post})\mathbf{1}_{\{w\in[21,30]\}} - \eta_L \cdot L_{wc}\end{aligned}$$

The differences $(\alpha_{city} - \eta_{city})$ and $(\alpha_{post} - \eta_{post})$ will depend on the amount of variance explained in T_{wc} that is shared between L_{wc} and $\mathbf{1}_{\{c=\text{DFW}\}}$ and between L_{wc} and $\mathbf{1}_{\{w\in[21,30]\}}$ respectively. The magnitude of η_L will depend on the amount of covariance between T_{wc} and $\tilde{L}_{wc}$.

Variance of $\hat{\gamma}_{treat}$:

Let us now consider what will happen to the variance on the treatment coefficient. Asymptotic normality of the OLS regression gives us the following expression for the estimated coefficient in (SA) (Wooldridge (2010)):

$$\sqrt{n}(\hat{\beta}_{treat} - \beta_{treat}) \xrightarrow{d} \mathcal{N}\left(0, \frac{\mathbb{E}[\varepsilon^2_{wc} | T_{wc}, \mathbf{1}_{\{c=\text{DFW}\}}, \mathbf{1}_{\{w\in[21,30]\}}]}{Var(\tilde{T}^A_{wc})}\right)$$

Given the assumptions on the residual term, we can simplify this to:

$$\sqrt{n}(\hat{\beta}_{treat} - \beta_{treat}) \xrightarrow{d} \mathcal{N}\left(0, \frac{(1 - R^2_{(SA)})}{Var(\tilde{T}^A_{wc})}\right)$$

Next the asymptotic distribution for the treatment coefficient in (SB) is given by:

$$\begin{aligned}\sqrt{n}(\hat{\gamma}_{treat} - \gamma_{treat}) &\xrightarrow{d} \mathcal{N}\left(0, \frac{\mathbb{E}[\mu^2_{wc} | T_{wc}, \mathbf{1}_{\{c=\text{DFW}\}}, \mathbf{1}_{\{w\in[21,30]\}}, L_{wc}]}{Var(\tilde{T}^B_{wc})}\right) \\ &\xrightarrow{d} \mathcal{N}\left(0, \frac{(1 - R^2_{(SB)})}{Var(\tilde{T}^B_{wc})}\right)\end{aligned}$$

Taking the ratio of the variances of $\hat{\gamma}_{treat}$ and $\hat{\beta}_{treat}$, we get the following expression:

$$\frac{(1 - R^2_{(SB)})}{(1 - R^2_{(SA)})} \cdot \frac{Var(\tilde{T}^A_{wc})}{Var(\tilde{T}^B_{wc})}$$

[23] Note, here $\tilde{L}_{wc}$ is the residualized version of L_{wc} based on the covariates in the specification of column 5.

Table 5 Auxiliary outcome metrics, DFW vs HOU

		Opportunities receiving						Equitable Proportion		
		$=1$ sign-up		$\geqslant 2$ sign-ups		$\geqslant 3$ sign-ups		Overall	Urban	Rural
City	Period	Mean	Std.	Mean	Std.	Mean	Std.			
Treated	Pre-exp.	137.0	19.3	45.7	16.4	18.0	9.1	0.557	0.549	0.565
(DFW)	Experiment	165.3	17.6	56.6	15.9	17.5	7.4	0.554	0.557	0.550
Control	Pre-exp.	95.6	11.8	22.7	6.4	6.5	2.8	0.663	0.676	0.644
(HOU)	Experiment	102.4	12.7	30.4	6.9	8.7	2.7	0.589	0.598	0.575
Change in the treated city (DFW)		+20.7%		+23.9%		-2.7%		-0.4%	+1.6%	-2.6%
Change in the control city (HOU)		+7.1%		+33.9%		+33.8%		-11.3%	-11.5%	-10.7%
Net change in DFW over HOU		**+13.6%**		**-10.0%**		**-36.5%**		**+10.9%**	**+13.1%**	**+8.1%**

Notice, when the correlation between $\log(Y_{wc})$ and L_{wc} (i.e., $\rho_{\log(Y_{wc}),L}$) is high, the first term of the above expression will be considerably less than 1. When the correlation between T_{wc} and $\tilde{L}_{wc}$ (i.e., $\rho_{T,\tilde{L}}$) is low, we have already shown that the second term of the above expression will be approximately 1. Thus, in this setting the variance of $\hat{\gamma}_{treat}$ is lower than that of $\hat{\beta}_{treat}$ asymptotically.

In our setting, the correlation between $\log(Y_{wc})$ and L_{wc} is $\rho_{\log(Y_{wc}),L} = 0.91$ while the correlation between T_{wc} and $\tilde{L}_{wc}$ is $\rho_{T,\tilde{L}} = -0.01$. This suggests that after controlling for L_{wc} (i.e., the log-transformed efficiency), we should see a decrease in the variance of our treatment coefficient. Table 2 shows exactly this pattern — the error on the coefficient in column 4 (0.027) is smaller than the error on the coefficient in column 3 (0.097) — suggesting that controlling for L_{wc} is appropriate. Again, the same argument holds for the specification in column 5. For this specification, $\rho_{\log(Y_{wc}),L} = 0.91$ and $\rho_{T,\tilde{L}} = -0.013$. The error on the coefficient in column 5 (0.025) is lower than that in column 3, suggesting that controlling for L_{wc} is appropriate here as well.

C.4. Distributional and Geographic Impacts of SmartSort

We now conduct a deeper investigation into the mechanism by which SmartSort achieves gains in equity. We first focus on the distributional impact of SmartSort before turning our attention to geographical differences in its impact.

Distributional Impacts of SmartSort. To understand the impact of SmartSort on equity beyond our primary outcome metric (defined in Section 4.1), here we consider the following auxiliary outcome metrics:

1. The weekly number of opportunities with *exactly* one sign-up.
2. The weekly number of opportunities with *at least* two sign-ups.
3. The weekly number of opportunities with *at least* three sign-ups.
4. The overall proportion of sign-ups that are "equitable," where a sign-up is defined as equitable if and only if the opportunity receiving the sign-up has *not* received a sign-up in the prior 7 days. Hereafter we refer to this metric as the *equitable proportion*.

Following the same steps as described in Section 4.2.1, in Table 5, we provide descriptive evidence on the impact of SmartSort on these auxiliary outcome metrics. The order of columns matches the order in which we define these metrics above.

Focusing for now on the first three columns of Table 5, these aggregate-level analyses suggest that the equity gain is primarily driven by an increase in the number of opportunities that receive exactly one sign-up. During the experiment period, the weekly number of opportunities with exactly one sign-up increases in both DFW and HOU (resp. 20.7% and 7.1%), however the increase is substantially higher in DFW (a net change of 13.6%). On the other hand, the net change in the weekly number of opportunities receiving 2 or more (resp. 3) is negative −10% (resp. −36.5%). The greater effect on the weekly number of opportunities receiving 3 or more sign-ups is particularly interesting, as it suggests that SmartSort does not just redistribute sign-ups away from opportunities at the margin.

These results are aligned with our expectation: because VM is under-supplied, every week there are many opportunities with no sign-ups. Among opportunities with similar levels of appeal, SmartSort will prioritize those with no sign-ups, thus increasing the number of opportunities getting a sign-up at the expense of those that would have gotten more sign-ups under CP.

The fourth column of Table 5 focuses on an alternative equity metric: the equitable proportion. To calculate this metric, we label each sign-up as equitable or not, depending on whether the corresponding opportunity had received a sign-up in the prior seven days. For each region (DFW and HOU) and each period (pre-experiment and experiment), we then compute the proportion of sign-ups which are equitable. Because SmartSort penalizes opportunities for seven days after receiving a sign-up (as described in Section 3.2), this metric more directly captures SmartSort's impact. Consistent with this intuition, we see a slightly larger net change in this equity metric (+10.9%) compared to the net change in our primary equity metric (+9.2%, see Table 1).

We complement this aggregate-level analysis with a formal DID analysis, with results presented in Table 6. The first three columns are concerned with the distributional impacts of SmartSort, and they use the following specification:

$$\log(Y_{wc}) = \beta_0 + \beta_{\text{treat}} \cdot \mathbf{1}_{\{w\in[21,30]\}} \cdot \mathbf{1}_{\{c=\text{DFW}\}} + \beta_{\text{city}} \cdot \mathbf{1}_{\{c=\text{HOU}\}} + \delta_w + \beta_S \cdot \log(S_{wc}) + \varepsilon_{wc}, \tag{72}$$

This specification is identical to the one used in column 5 of Table 2; it controls for city and week fixed effects as well as the log-transformed efficiency, denoted by $\log(S_{wc})$. We observe that the estimated treatment effects are directionally aligned with our aggregate-level analysis in Table 5.

The fourth column of Table 6 is concerned with the equitable proportion. This analysis is conducted on the sign-up level, and hence the number of observations is much larger. We use the following specification:

$$Y_{swc} = \beta_0 + \beta_{\text{treat}} \cdot \mathbf{1}_{\{w\in[21,30]\}} \cdot \mathbf{1}_{\{c=\text{DFW}\}} + \beta_{\text{city}} \cdot \mathbf{1}_{\{c=\text{HOU}\}} + \delta_w + \varepsilon_{swc}, \tag{73}$$

where Y_{swc} is an indicator which equals 1 if sign-up s in week w and city c went to an opportunity that had not received a sign-up in the past week. Roughly speaking, the estimate for β_{treat} can be interpreted as a 7.6 percentage point increase in the probability that a randomly chosen sign-up from DFW during the experiment period is equitable, i.e., went to an opportunity that had not received a sign-up in the past week. This effect size is quite large compared to the baseline equitable proportion of 0.557.

Table 6 DID Analysis of Auxiliary Outcome Metrics, DFW vs HOU

	Number of Opportunities Receiving				
	$=1$ sign-up	$\geqslant 2$ sign-ups	$\geqslant 3$ sign-ups	Equitable Proportion	
Treated	**0.125**	**-0.060**	**-0.288**	**0.076*****	**0.087*****
	(0.067)	(0.089)	(0.233)	(0.021)	(0.025)
Treated × Rural					**-0.022**
					(0.026)
Rural					0.017
					(0.019)
City	-0.132	-0.147	0.683	0.107***	0.126***
	(0.144)	(0.196)	(0.342)	(0.016)	(0.019)
City × Rural					-0.045
					(0.026)
Log Sign-ups	0.425	0.999**	3.116***		
	(0.253)	(0.319)	(0.519)		
Intercept	2.491	-1.772	-14.337***	0.526***	0.517***
	(1.418)	(1.791)	(2.901)	(0.025)	(0.027)
Week FE	Yes	Yes	Yes	Yes	Yes
Obs.	40	40	40	9015	9015
R^2	0.969	0.982	0.925	0.014	0.15
Adj. R^2	0.929	0.959	0.828	0.012	0.12
RSE	0.068	0.092	0.250	0.491	0.491
F Stat.	44.84***	65.54***	9.18***	6.31***	5.68***

Notes: Robust standard errors are reported in parentheses.
*p<0.05; **p<0.01; ***p<0.001

Geographic Impacts of SmartSort. Here we use data from our experiment in DFW to investigate heterogeneity in SmartSort's impact across geographic areas. Based on the simulations in Figure 4 and the discussion in Appendix B.7, we would expect that the gains in equity would be larger in areas where there is a greater density of high-appeal opportunities. One natural way to test this hypothesis is to divide each city into two groups (consisting of urban and rural zip codes) and repeat the analysis of Section 4.2 for each group separately. Specifically, to approximately balance the two groups, we define a zip code as urban if its population density exceeds 3000 people per square mile; otherwise, we define the zip code as rural. Under this definition, 52% of the sign-ups in DFW during the pre-experiment period went to opportunities in urban zip codes, even though these zip codes account for only 10% of the treated area. In HOU, 61% of sign-ups during the pre-experiment period went to opportunities in urban zip codes, and these zip codes account for only 15% of the control area.[24]

In Table 7, we present descriptive statistics for the urban and rural areas separately. This is the analog of Table 1, but for two separate subsets of the data. We observe that the net change in equity is notably larger in the urban zip codes than the rural zip codes (by 8.8 percentage points), which is consistent with our

[24] There are some zip codes (accounting for less than 5% of sign-ups) which do not have an associated area, e.g., due to being a collection of P.O. Boxes. We follow the convention of counting all such zip codes as urban zip codes.

expectations. The net change in efficiency is slightly positive in the urban group and moderately negative in the rural group. This overall difference (of 5.2 percentage points) may reflect the larger efficiency cost of imposing a penalty in areas with fewer high-appeal opportunities, as illustrated by our simulation study. It could also simply be a consequence of some relative shifts in sign-ups between urban and rural zip codes due to the substantial amount of activity on the border between the two. (This is in contrast to the overall treated area, which has limited activity along its boundary; see the discussion in Section 4.1.1.)

Table 7 Platform's city-wide weekly average counts, DFW vs HOU, separated into urban and rural groups

Urban		Efficiency		Equity	
City	Period	Mean	Std.	Mean	Std.
Treated (DFW)	Pre-exp.	136.5	42.8	94.4	20.9
	Exp.	160.9	37.5	116.0	16.1
Control (HOU)	Pre-exp.	92.0	19.9	72.7	13.2
	Exp.	107.5	18.7	80.0	16.4
Change in DFW		+17.9%		+22.9%	
Change in HOU		+16.8%		+10.0%	
Net change		+1.1%		+12.9%	

Rural		Efficiency		Equity	
City	Period	Mean	Std.	Mean	Std.
Treated (DFW)	Pre-exp.	125.1	28.4	88.3	14.8
	Exp.	147.9	33.5	105.9	18.7
Control (HOU)	Pre-exp.	59.2	9.6	45.6	6.8
	Exp.	72.4	9.6	52.8	7.9
Change in DFW		+18.2%		+19.9%	
Change in HOU		+22.3%		+15.8%	
Net change		-4.1%		+4.1%	

As an alternative way of comparing equity in the urban and rural groups, we turn to one of our auxiliary outcome metrics: the equitable proportion. In the last two columns of Table 5, we report the average for this metric for the urban and rural groups separately. Consistent with our expectations, we observe a larger relative increase in the equitable proportion in the urban area (a net change of 13.1%) compared to the rural area (a net change of 8.1%).

We present a formal DID analysis in Table 6. Specifically, in Column 5 we include information about the sign-up's zip code z in an effort to identify any heterogeneity in SmartSort's geographic impacts between urban and rural locations. We use the following specification, which introduces an interaction between the treatment indicator and an indicator for being in the set of rural zip codes $\mathcal{R}$.

$$\begin{aligned} Y_{swcz} \;=\; & \beta_0 + \beta_{\text{treat}} \cdot \mathbf{1}_{\{w\in[21,30]\}} \cdot \mathbf{1}_{\{c=\text{DFW}\}} + \beta_{\text{city}} \cdot \mathbf{1}_{\{c=\text{HOU}\}} + \beta_{\text{rural}} \cdot \mathbf{1}_{\{z\in\mathcal{R}\}} \\ & + \beta_{\text{city}\times\text{rural}} \cdot \mathbf{1}_{\{c=\text{HOU}\}} \cdot \mathbf{1}_{\{z\in\mathcal{R}\}} + \beta_{\text{treat}\times\text{rural}} \cdot \mathbf{1}_{\{z\in\mathcal{R}\}} \cdot \mathbf{1}_{\{w\in[21,30]\}} \cdot \mathbf{1}_{\{c=\text{DFW}\}} + \delta_w + \varepsilon_{swcz}, \end{aligned} \tag{74}$$

The estimate for β_{treat} suggests an increase of 0.087 in the equitable proportion during the experiment period *for sign-ups in DFW's urban zip codes*. We estimate a smaller increase for DFW's rural zip codes: when combining the coefficient β_{treat} with the coefficient on the interaction term $\beta_{\text{treat}\times\text{rural}}$, our results imply an increase of only 0.065 in the equitable proportion in rural zip codes. This geographic heterogeneity – a larger equity gain in dense urban areas – is directionally consistent with our simulation results in Figure 4 and the thought experiment described in Appendix B.7, though we note that this difference is not statistically significant.

Appendix D: Empirical Analysis for the Southern California Experiment

D.1. Results in Southern California

Upon completion of the first wave of experimentation in DFW, we designed an additional wave of experimentation in three California regions. This wave of experimentation followed a staggered roll-out, such that each region adopted SmartSort at a different time. The order of roll-out was chosen at random.

The timeline for the roll-out in California is as follows:

1. Phase I began on Saturday, November 5, 2022 in West LA, defined as all zip codes with a geographic center within the California state boundaries with a latitude between 34.5° N and 33.6° N and longitude west of 118.1° W.
2. Phase II began on Tuesday, December 6, 2022 in East LA, defined as all zip codes with a geographic center within the California state boundaries with a latitude between 34.5° N and 33.6° N and longitude east of 118.1° W.
3. Phase III began on Wednesday, January 4, 2023 in San Diego, defined as all zip codes with a geographic center within the California state boundaries with a latitude south of 33.6° N.

We emphasize that once the treatment was turned on in each of these regions, it was not turned off.

Given the variability in the time of treatment adoption for these geographic regions, we would ideally like to run a staggered treatment adoption analysis to obtain our difference-in-differences estimates. However, these three regions have contiguous borders (see Figure 3). Opportunities near these borders could be viewed by both volunteers that are experiencing SmartSort search rankings and those that are experiencing search rankings based on CP. This phenomenon could create contamination that would violate the no anticipation assumption required for a difference-in-differences analysis. For example, the introduction of SmartSort in West LA might create an effect in East LA *prior* to the introduction of SmartSort in East LA. To avoid this concern, we first consider an identical analysis to that in DFW in which we combine all California experimental regions into a single Southern California region (SCA) and we compare the last ten weeks before Phase I (i.e., the last ten weeks when no part of the region was treated) to the first ten weeks after Phase II (i.e., the first ten weeks when the entire region was treated).[25] Specifically we define the pre-experiment period to be Tuesday August 23, 2022 to Monday October 31, 2022 (corresponding to weeks 34-43 of 2022) and the experiment period to be Wednesday January 4, 2023 to Tuesday March 14, 2023 (corresponding to weeks 1-10 of 2023).[26]

To control for seasonal trends, we compare SCA to the San Francisco/San Jose region (denoted by SFJ), which consists of all zip codes that are within a 20-mile radius of either the San Francisco or San Jose city centers (measured using Haversine distance). This control region is not contiguous with the treated region.

[25] We emphasize that conducting a separate analysis for each phase leads to similar results; however, we present this aggregate analysis to avoid any concerns about interference across phases.

[26] In keeping with the implementation schedule in our first experiment, we had intended the final California roll-out to be on Tuesday January 3, 2023. However, due to an administrative issue at VolunteerMatch the actual implementation date was Wednesday January 4, 2023, resulting in experiment period weeks that run from Wednesday to Tuesday of each week. An identical analysis that defines the pre-experiment period as Wednesday August 24, 2022 to Tuesday November 1, 2022 produces nearly identical results.

SCA and SFJ each contain the two largest cities in their area, are comparable over a spectrum of observable characteristics (see Table 11 in Appendix D.2), and experience similar seasonal trends (see Figure 11 in Appendix D.2). To be consistent with the terminology of our first experiment, we refer to SCA and SFJ as "city," even though they are comprised of more than one city. The data we will be analyzing in the second experiment is based on sign-ups in both the pre-experiment and experiment periods for opportunities located within the treated area of SCA and the control area of SFJ. To avoid overly influential outliers, we drop the top 1% of opportunities in each of the treated and control regions by sign-ups received. The resulting dataset, which we will use throughout this section, consists of sign-ups for 5,886 unique opportunities in SCA and 2,298 unique opportunities in SFJ.

Aggregate statistics from this second experiment show a similar effect as we saw in our first experiment. Table 8 — which is the analog of Table 1 from Section 4.2.1 — shows that both SCA and SFJ experience an increase in efficiency (presumably due to an increase in volunteering activity at the beginning of the year); however the net change in SCA over SFJ is quite small in magnitude (−0.9%). On the other hand, while both SCA and SFJ also experience an increase in equity, the net change in SCA over SFJ is large and positive (+8.9%).

Table 8 Platform's city-wide weekly average counts, SCA vs SFJ

		Efficiency		Equity	
City	Period	Mean	Std.	Mean	Std.
Treated (SCA)	Pre-experiment	1258.9	91.9	851.3	57.8
	Experiment	1503.9	127.7	1035.0	59.4
Control (SFJ)	Pre-experiment	544.7	69.8	358.6	36.9
	Experiment	655.8	82.9	404.2	38.2
Change in the treated region (SCA)		+19.5%		+21.6%	
Change in the control region (SFJ)		+20.4%		+12.7%	
Net change in SCA over SFJ		-0.9%		**+8.9%**	

The full difference-in-differences analysis supports these aggregate results. Table 9 (the analog of Table 2 in Section 4.2.2) summarizes this analysis. The coefficient of -0.010 in columns 1 and 2 can be interpreted as a small, statistically insignificant 1.0% decrease in the total number of sign-ups in a week after SmartSort was introduced in 2023 in SCA.

Meanwhile, we do see significant changes in equity; the treatment coefficients (0.075, 0.082, and 0.083, respectively) in columns 3, 4, and 5 correspond to an increase (of 7.8%, 8.5%, and 8.7%, respectively) in equity. In columns 4 and 5 (i.e., after controlling for log-transformed efficiency), we see that the increase in equity (given a fixed level of efficiency) is statistically significant.[27]

[27] The conditions stated in Appendix C.3 that motivate controlling for (log-transformed) efficiency in Columns 4 and 5 of Table 9 also hold in our second setting. Details are omitted for the sake of brevity.

Table 9 Log Transformed City Level Analyses, SCA vs SFJ

	Efficiency		Equity		
	(1)	(2)	(3)	(4)	(5)
Treated	-0.010	-0.010	0.075	0.082***	0.083**
	(0.071)	(0.065)	(0.056)	(0.019)	(0.024)
Post	0.187**		0.121*	-0.017	
	(0.061)		(0.047)	(0.016)	
City	0.843***	0.843***	0.867***	0.248***	0.232
	(0.052)	(0.045)	(0.042)	(0.042)	(0.127)
Log Sign-ups				0.734***	0.754***
				(0.045)	(0.155)
Intercept	6.292***	6.303***	5.877***	1.258***	1.160
	(0.045)	(0.055)	(0.035)	(0.282)	(0.979)
Week FE	No	Yes	No	No	Yes
Obs.	40	40	40	40	40
R^2	0.947	0.989	0.971	0.997	0.999
Adj. R^2	0.942	0.976	0.969	0.997	0.997
RSE	0.107	0.069	0.083	0.028	0.024
F Stat.	188.8***	86.9***	387.9***	2325.6***	358.6***

Notes: Robust standard errors are reported in parentheses.
*p<0.05; **p<0.01; ***p<0.001

Table 10 Auxiliary outcome metrics, SCA vs SFJ

		Opportunities receiving						Equitable Proportion		
		$=1$ sign-up		$\geqslant 2$ sign-ups		$\geqslant 3$ sign-ups		Overall	Urban	Rural
City	Period	Mean	Std.	Mean	Std.	Mean	Std.			
Treated	Pre-exp.	613.1	40.1	238.2	27.8	85.5	13.4	0.508	0.498	0.545
(SCA)	Experiment	744.4	42.9	290.6	33.0	91.6	24.2	0.538	0.528	0.573
Control	Pre-exp.	251.6	21.7	107.0	16.7	41.0	9.3	0.476	0.477	0.471
(SFJ)	Experiment	263.5	24.0	140.7	19.5	58.6	12.6	0.438	0.440	0.428
Change in the treated city (SCA)		+21.4%		+22.0%		+7.1%		+5.9%	+6.2%	+5.1%
Change in the control city (SFJ)		+4.7%		+31.5%		+42.9%		-8.0%	-7.8%	-9.1%
Net change in SCA over SFJ		**+16.7%**		**-9.5%**		**-35.8%**		**+13.9%**	**+13.9%**	**+14.2%**

Digging deeper into the distributional impacts of SmartSort in SCA, we see that here too the equity gain was primarily driven by an increase in the number of opportunities receiving exactly one sign-up. Table 10 (the analog of Table 5) presents the descriptive statistics for the outcome metrics described in Appendix C.4. In the left three columns, we observe that during the experiment period, the weekly number of opportunities with exactly one sign-up increases in both SCA and SFJ (resp. 21.4% and 4.7%), however the increase is substantially higher in SCA (a net change of +16.7%). On the other hand, the net change in the weekly

number of opportunities receiving at least 2 (resp. 3) sign-ups is −9.5% (resp. −35.8%). This pattern of redistribution is again aligned with our expectation.

In the fourth column of Table 10, we present the equitable proportion for the entire area, followed by a breakdown into urban and rural zip codes. The net change in the overall equitable proportion is consistent with our findings in DFW, as we observe an even greater change in this equity metric (which we remind more directly captures SmartSort's impact). We observe minimal geographic heterogeneity, possibly due to the higher density in these two regions. In SCA and SFJ, more than 80% of the sign-ups during the pre-experiment period occurred in zip codes that meet our definition of urban – at least 3000 people per square mile – compared to around 55% of the sign-ups in DFW and HOU.

D.2. Supplemental Evidence for the Southern California Experiment

SCA and SFJ are two of the largest regions in California. This makes SFJ a natural comparison city. Drawing on zip-level data from the US Census American Community Survey 2020,[28] in Table 11 we show that in terms of population, demographic composition, and income level the regions included in our analysis are fairly comparable though SCA is much larger. Additionally, using aggregate city-level data on volunteer activity provided by Google Analytics we also confirm that in our pre-experiment period (defined in Appendix D.1) the number of visitors and sessions (i.e., the periods of time during which visitors interact with the VM site) in the two largest cities in each region are on similar orders; we note that this data is not available on a more granular level, and it comprises a subset of the full activity in the SCA and SFJ regions.

Table 11 Balance Table - SCA and SFJ

	SCA	SFJ
Total Population	20,362,390	3,380,356
Pct Population Female	50.48%	50.17%
Pct Population White	53.44%	45.52%
Median Age (years)	36.6	38.2
Median Income	$78,000	$120,000
Total Users (pre-experiment)	40,903	18,109
Total Sessions (pre-experiment)	58,218	26,049

The key identifying assumption for the DID analyses is that of parallel trends between the treatment and control groups. Here we provide evidence in support of the parallel trend assumption between SCA and SFJ. In Figure 11 we simply plot the trends for our two outcome metrics across the two regions. Recall that the pre-experiment ran for 10 consecutive weeks from Tuesday August 23, 2022 - Monday October 31, 2022 (week 34 - week 43) and the experiment period ran for 10 consecutive weeks starting from Wednesday January 4, 2023 - Tuesday March 14, 2023 (week 1 - week 10). For both metrics, we observe strong and

[28] https://www.census.gov/newsroom/press-kits/2021/acs-5-year.html

reasonably consistent seasonal trends in the pre-experiment period, with a noticeable start-of-the-year boost in volunteer activity in both California regions. Further evidence of the parallel trends assumption is available but omitted for brevity.

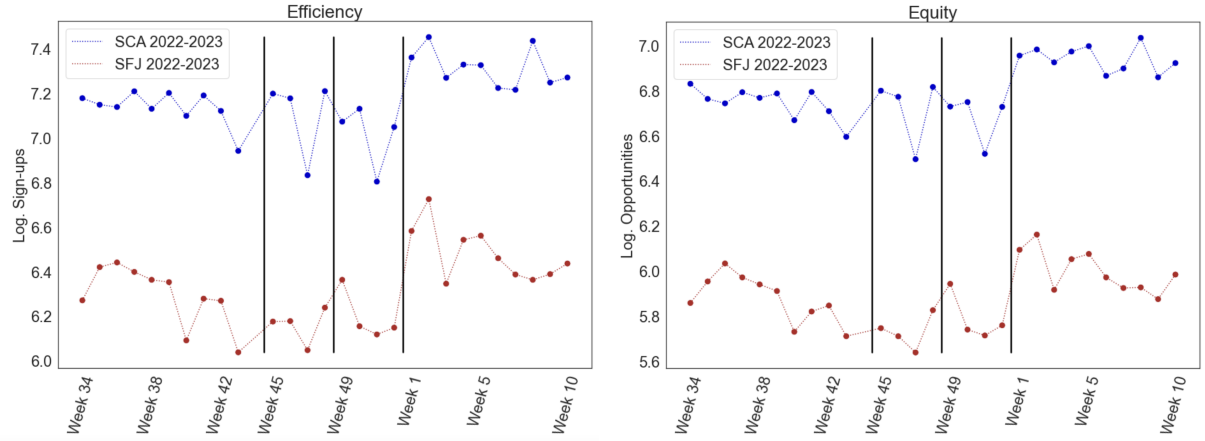

(a) (b)

Figure 11 The weekly seasonal trends of efficiency and equity in SCA and SFJ.