
Fair Sortition Made Transparent

Bailey Flanigan, Gregory Kehne, and Ariel D. Procaccia

Abstract

Sortition is an age-old democratic paradigm, widely manifested today through the random selection of citizens’ assemblies. Recently-deployed algorithms select assemblies *maximally fairly*, meaning that subject to demographic quotas, they give all potential participants as equal a chance as possible of being chosen. While these fairness gains can bolster the legitimacy of citizens’ assemblies and facilitate their uptake, existing algorithms remain limited by their lack of transparency. To overcome this hurdle, in this work we focus on panel selection by uniform lottery, which is easy to realize in an observable way. By this approach, the final assembly is selected by uniformly sampling some pre-selected set of m possible assemblies. We provide theoretical guarantees on the fairness attainable via this type of uniform lottery, as compared to the existing maximally fair but opaque algorithms, for two different fairness objectives. We complement these results with experiments on real-world instances that demonstrate the viability of the uniform lottery approach as a method of selecting assemblies both fairly and transparently.

1 Introduction

In a *Citizens’ Assembly*, a panel of randomly-chosen citizens is convened to deliberate and ultimately make a recommendation on a policy issue. The defining aspect of citizens’ assemblies is the randomness of the process, *sortition*, by which their participants are chosen. In practice, the sortition process works as follows: first, volunteers are solicited via thousands of letters or phone calls, which target households chosen uniformly at random. Those who respond affirmatively form the *pool* of volunteers, from which the final panel will be chosen. Then, a *selection algorithm* is used to randomly select some pre-specified number k of pool members for the panel. To ensure adequate representation of demographic groups, the chosen panel must satisfy some upper and lower quotas on orthogonal feature categories, such as age, gender, and ethnicity. We call such a panel a *feasible panel*. As this process illustrates, citizens’ assemblies offer a way to involve the public in informed decision-making. As a result of this unique property, over the past decades, citizens’ assemblies have been gaining momentum globally, even being commissioned by governments and leading to policy changes at the national level [ICAP19, OEC20, FGG⁺21].

Prompted by the growing impact of citizens’ assemblies, there has been a recent flurry of computer scientific research on sortition, and in particular, on the fairness of the procedure by which participants are chosen [BGP19, FGGP20, FGG⁺21]. The most practicable result to date is a family of selection algorithms proposed by Flanigan *et al.* [FGG⁺21], which are distinguished from their predecessors by their use of randomness toward the goal of fairness: while previously-used algorithms selected pool members in a random but ad-hoc fashion, these new algorithms are *maximally fair*, ensuring that pool members have as equal probability as possible of being chosen for the panel, subject to the quotas.¹ To encompass the many interpretations of “as equal as possible,” these algorithms permit the optimization of any fairness objective with certain convexity properties. There is now a publicly available implementation of Flanigan *et al.*’s techniques, called *Panelot*, which optimizes the egalitarian notion that no pool member has too little selection probability via the *Leximin* objec-

¹Quotas can preclude giving individuals exactly equal probabilities: if the panel must be 1/2 men, 1/2 women but the pool is split 3/4 men, 1/4 women, then some women must be chosen more often than some men.

tive from Fair Division [Mou03, FGG⁺20]. This algorithm has already been deployed by several groups of panel organizers, and has been used to select dozens of panels worldwide.

Fairness gains in the panel selection process can lend legitimacy to citizens' assemblies and potentially increase their adoption, but only insofar as the public trusts that these gains are truly being realized. Currently, the potential for public trust in the panel selection process is limited by multiple factors. First, even the latest panel selection algorithms are most conducive to selecting the final panel via behind-the-scenes computation. When panels are selected in this manner, people cannot even verify that every pool member has *any* chance of being chosen for the panel. A second and more fundamental hurdle is that randomness and probability, which are so central to the sortition process, have been shown in many contexts to be difficult for people to understand and reason about [RB08, MP16, WML20]. Aiming to address these shortcomings, we define and pursue the following notion of *transparency* in panel selection:

Transparency: Onlookers should be able to, without reasoning in-depth about probability, (1) understand the probabilities with which each individual will be chosen for the panel *in theory*, and (2) verify that individuals are actually selected with these probabilities *in practice*.

In this paper, we aim to achieve transparency and fairness simultaneously: this means advancing the defined goal of transparency, while preserving the fairness gains introduced by maximally fair selection algorithms. Although this task is reminiscent of existing AI research on trade-offs between fairness or transparency with other desirable objectives [BFT12, FKL16, BHJ⁺18, TBG⁺20], to our knowledge, this is one of the first investigations of the trade-off between fairness and transparency.

Setting aside for a moment the goal of fairness, we consider a method of random decision-making that is already common in the public sphere: the uniform lottery. To satisfy quotas, a uniform lottery for sortition must randomize not over individuals, but over entire feasible panels. In fact, this approach has been suggested by practitioners, and was even used in 2020 to select a citizens' assembly in Michigan. The following example, which closely mirrors that real-world pilot,² illustrates that panel selection via uniform lottery is naturally consistent with the transparency notion we pursue.

Suppose we construct 1000 feasible panels from a pool (possibly with duplicates), numbered 000-999, and publish an (anonymized) list of which pool members are on each panel. We then inform spectators that we will choose each panel with equal probability. This satisfies criterion (1): spectators can easily understand that all panels will be chosen with the same probability of 1/1000, and with only slightly more reasoning, they can observe each individual's selection probabilities by counting the number of panels containing them. To satisfy criterion (2), we enact the lottery by drawing each of the 3 digits of the final panel number individually from lottery machines. Lottery spectators can watch each ball be drawn with equal probability, providing confirmation that panels are indeed being chosen with uniform probabilities (and thus confirming the enactment of the proposed individual selection probabilities). In addition to be mainstream, decision-making via drawing lottery balls invites an exciting spectacle, which can promote engagement with citizens' assemblies.

This simple method neatly satisfies our transparency criteria, but it has one obvious downside: a uniform lottery over an arbitrary set of feasible panels does not guarantee any measure of equal probabilities to individuals. In fact, it is not even clear that the *fairest possible* uniform lottery over m panels, where m is a number conducive to selection by physical lottery (e.g. $m = 1000$), would not be significantly less fair than maximally fair algorithms. For example, if m is too small, there may be *no* uniform lottery which gives all individuals non-zero selection probability, even if each individual appears on some feasible panel.

Fortunately, empirical evidence suggests that there is hope: in the 2020 pilot mentioned above, a uniform lottery over $m = 1000$ panels was found that nearly matched the fairness of the maximally fair distribution generated by Panelot. Motivated by this anecdotal evidence, we aim to understand whether such a fair uniform lottery is guaranteed to exist in general, and if it does, how to find it. We summarize this goal in the following research questions:

²Of By For's pilot of live panel selection via lottery can be viewed at <https://vimeo.com/458304880#t=17m59s> from 17:59 to 21:23. For a more detailed description, see Figure 3 and surrounding text in [FGG⁺21].

Does there exist a uniform lottery over m panels that nearly preserves the fairness of the maximally fair unconstrained distribution over panels? And, Algorithmically, how do we compute such a uniform lottery?

Results and contributions. After describing the model in Section 2, in Section 3 we prove that it is possible to round an (essentially) arbitrary distribution over panels to a uniform lottery while preserving *all* individuals’ selection probabilities up to only a small bounded deviation. These results use tools from discrepancy theory and randomized rounding. Intuitively, this bounded change in selection probabilities implies bounded losses in fairness; we formalize this intuition in Section 4, showing that there exists in general a uniform lottery that is nearly maximally fair, with respect to multiple fairness objectives. Although we would ideally like to give such bounds for the Leximin fairness objective due to its use in practice, this objective cannot be summarized by a single expression. Thus, we give bounds for the closely-related egalitarian objective, *Maximin* [CELM07]. We additionally give upper bounds on the loss in *Nash Welfare* [Mou03], a similarly well-established fairness objective that has also been implemented in panel selection tools [HG20].

Finally, in Section 5, we consider the algorithmic question: given a maximally fair distribution over panels, can we *find* a near-maximally fair uniform lottery, as our bounds suggest should exist? To answer this question, we implement two standard rounding algorithms, along with near-optimal (but more computationally-intensive) integer programming methods, for finding uniform lotteries. We then evaluate the performance of these algorithms in 11 real-world panel selection instances. We find that in all instances, we can compute uniform lotteries that nearly exactly preserve not only fairness with respect to both objectives, but *entire sets* of Leximin-optimal marginals, meaning that from the perspective of individuals, there is essentially no difference between using a uniform lottery versus the optimal distribution used by the latest algorithms. We discuss these results, their implications, and how they can be deployed directly into the existing panel selection pipeline in Section 6.

2 Model

Panel Selection Problem. First, we formally define the task of panel selection for citizens’ assemblies. Let $N = [n]$ be the *pool* of volunteers for the panel—individuals from the population who have indicated their willingness to participate in response to an invitation. Let $F = \{f_t\}_t$ denote a fixed set of *features* of interest. Each feature $f_t : N \rightarrow \Omega_t$ maps each pool member to their value of that feature, where Ω_t is the set of f_t ’s possible values. For example, for feature $f_t = \text{“gender”}$, we might have $\Omega_t = \{\text{“male”}, \text{“female”}, \text{“non-binary”}\}$. We define individual i ’s *feature vector* $F(i) = (f_t(i))_t \in \prod_t \Omega_t$ to be the vector encoding their values for all features in F .

As is done in practice and in previous research [FGGP20, FGG⁺21], we impose that the chosen panel P must be a subset of the pool of size k , and must be representative of the broader population with respect to the features in F . This representativeness is imposed via *quotas*: for each feature f and corresponding value $v \in \Omega$, we may have lower and upper quotas $l_{f,v}$ and $u_{f,v}$. These quotas require that the panel contain between $l_{f,v}$ and $u_{f,v}$ individuals i such that $f(i) = v$.

In terms of these parameters, we define an instance of the panel selection problem as: given (N, k, F, l, u) —a pool, panel size, set of features, and sets of lower and upper quotas—randomly select a *feasible panel*, where a feasible panel is any set of individuals P from the collection \mathcal{K} :

$$\mathcal{K} := \{P \in \binom{N}{k} : l_{f,v} \leq |\{i \in P : f(i) = v\}| \leq u_{f,v} \text{ for all } f, v\}.$$

Maximally Fair Selection Algorithms. A *selection algorithm* is a procedure that solves instances of the panel selection problem. A selection algorithm’s level of fairness on a given instance is determined by its *panel distribution* p , the (possibly implicit) distribution over \mathcal{K} from which it draws the final panel. Because we care about fairness to individuals, we evaluate the fairness of p in terms of the individual selection probabilities, or *marginals*, that p implies.³ We denote the vector of marginals implied by p as π , and we will sometimes specify a panel distribution as p, π to explicitly denote this pair. We say that π is *realizable* if it is implied by some valid panel distribution p .

Maximally fair selection algorithms are those which solve the panel selection problem by sampling a specifically chosen p : one which implies marginals π that allocate probability as fairly as possible

³Any distribution over panels p implies a selection probability for each pool members: A pool member’s selection probability, per p , is equal to the probability of drawing a panel from p containing that pool member.

across pool members. The fairness of p, π , is measured by a *fairness objective* \mathcal{F} , which maps an allocation—in this case, of selection probability to pool members—to a real number measuring the allocation’s fairness. Fixing an instance, a fairness objective \mathcal{F} , and a panel distribution p , we express the fairness of p as $\mathcal{F}(p)$. Existing maximally fair selection algorithms can maximize a wide range of fairness objectives, including the three we consider in this paper.

Leximin, Maximin, and Nash Welfare. Of the three fairness objectives we consider in this paper, Maximin and Nash Welfare (NW) have succinct formulae. For p, π they are defined as follows, where π_i is the marginal of individual i :

$$\text{Maximin}(p) := \min_{i \in N} \pi_i, \quad \text{NW}(p) := \left(\prod_i \pi_i \right)^{1/n}.$$

Intuitively, NW maximizes the geometric mean, prioritizing the marginal π_i of each individual i in proportion to π_i^{-1} . Maximin maximizes the marginal probability of the individual least likely to be selected. Finally, Leximin is a refinement of Maximin, and is defined by the following algorithm: first, optimize Maximin; then, fixing the minimum marginal as a lower bound on any marginal, maximize the second-lowest marginal; and so on.

Our setting. In this paper, we study m -uniform panel distributions \bar{p} —distributions over \mathcal{K} in which all probabilities are multiples of $1/m$. Formally, while an unconstrained panel distribution p lies in $\mathcal{D} := \{p \in \mathbb{R}_+^{|\mathcal{K}|} : \|p\|_1 = 1\}$, “uniform lotteries” \bar{p} for fixed m lie in $\bar{\mathcal{D}} := \{\bar{p} \in (\mathbb{Z}_+/m)^{|\mathcal{K}|} : \|\bar{p}\|_1 = 1\}$. In particular, for each aforementioned objective \mathcal{F} , we want to study the existence and construction of a uniform distribution \bar{p} that is nearly as fair as $p^* \in \arg \max_{p \in \mathcal{D}} \mathcal{F}(p)$, the maximally fair unconstrained distribution, which is sampled by existing maximally fair selection algorithms. To measure the extent to which such a lottery \bar{p} is achievable, we define the *fairness loss* between two panel distributions p, \bar{p} as $\mathcal{F}(p) - \mathcal{F}(\bar{p})$. We ultimately bound the fairness loss by first bounding the *marginal discrepancy*, which we define for $\pi, \bar{\pi}$ as $\|\pi - \bar{\pi}\|_\infty$.

3 Theoretical Bounds on Marginal Discrepancy

Here we prove that for a fixed panel distribution p, π , there exist uniform lotteries $\bar{p}, \bar{\pi}$ such that $\|\pi - \bar{\pi}\|_\infty$ is bounded. Preliminarily, we note that it is intuitive that bounds on this discrepancy should approach 0 as m becomes large with respect to n and k . To see why, begin by fixing some distribution p, π over panels: as m becomes large, we approach the scenario in which a uniform lottery \bar{p} can assign panels arbitrary probabilities, providing increasingly close approximations to p . Then, since the marginals π_i are continuous with respect to p , as $\bar{p} \rightarrow p$ we have that $\bar{\pi}_i \rightarrow \pi_i$ for all i .

While this argument demonstrates convergence, it provides neither efficient algorithms nor tight bounds on the rate of convergence. In this section, our task is therefore to bound the rate of this convergence as a function of m and the other parameters of the instance. All omitted proofs of results from this section are included in Appendix B.

3.1 Worst-Case Upper Bounds

Our first set of upper bounds result from rounding STANDARD LP, the LP that most naturally arises from our problem. This LP is defined in terms of a panel distribution p, π , and M , an $n \times |\mathcal{K}|$ matrix describing which individuals are on which feasible panels, i.e., $M_{i,P} = 1$ if $i \in P$ and $M_{i,P} = 0$ otherwise.

$$\text{STANDARD LP} \quad Mp = \pi \tag{3.1}$$

$$\begin{aligned} \|p\|_1 &= 1 \\ p &\geq 0. \end{aligned} \tag{3.2}$$

Here, (3.1) specifies n total constraints. Our goal is to round p to a uniform lottery \bar{p} over m panels (so the entries \bar{p} are multiples of $1/m$) such that (3.2) is maintained exactly, and no constraint in (3.1) is relaxed by too much, i.e., $\|Mp - M\bar{p}\|_\infty = \|\pi - \bar{\pi}\|_\infty$ remains small.

Randomized rounding is a natural first approach. Any randomized rounding scheme satisfying negative association (which includes several that respect (3.2)) yields the following bound:

Theorem 3.1. *For any realizable π , we may efficiently randomly generate \bar{p} such that its marginals $\bar{\pi}$ satisfy*

$$\|\pi - \bar{\pi}\|_\infty = O\left(\frac{\sqrt{n \log n}}{m}\right).$$

Fortunately, there is potential for improvement: randomized rounding does not make full use of the fact that M is k -column sparse, due to each panel in \mathcal{K} containing exactly k individuals. We use this sparsity to get a stronger bound for the practically significant parameter regime $n \gg k^2$. The proof applies a dependent rounding algorithm based on a theorem of Beck and Fiala [BF81], to which a modification ensures the exact satisfaction of constraint (3.2).

Theorem 3.2. *For any realizable π , we may efficiently construct \bar{p} such that its marginals $\bar{\pi}$ satisfy*

$$\|\pi - \bar{\pi}\|_\infty \leq k/m.$$

3.2 Beyond-Worst-Case Upper Bounds

As we will demonstrate in Section 3.3, we cannot hope for a better worst-case upper bound than $\text{poly}(k)/m$. We thus shift our consideration to instances which are “simple” in their feature structure, having a small number of features (Theorem B.5), a limited number of unique feature vectors in the pool (Theorem 3.3), or multiple individuals that share each feature vector present (Theorem B.6). The beyond-worst-case bounds given by Theorem 3.3 and Theorem B.6 asymptotically dominate our worst-case bounds in Theorem 3.1 and Theorem 3.2, respectively. Moreover, Theorem 3.3 dominates all other upper bounds in 10 of the 11 practical instances studied in Section 5.

We note that while our worst-case upper bounds implied the near-preservation of *any* realizable set of marginals π , some of our beyond-worst-case results apply to only realizable π which are *anonymous*, meaning that π_i are equal for all i with equal feature vectors. We contend that any reasonable set of marginals should have this property,⁴ and furthermore that the “anonymization” of any realizable π is also realizable (Claim B.4); hence this restriction is insignificant. Our beyond-worst-case bounds also differ from our worst-case bounds in that they depart from the paradigm of rounding p , instead randomizing over panels that may fall outside the support of p .

The main beyond-worst-case bound we give, stated below, is parameterized by $|\mathcal{C}|$, where \mathcal{C} is the set of unique feature vectors that appear in the pool. All omitted proofs and other beyond worst-case results are stated and proven in Appendix B.

Theorem 3.3. *If π is anonymous and realizable, then we may efficiently construct \bar{p} such that its marginals $\bar{\pi}$ satisfy*

$$\|\pi - \bar{\pi}\|_\infty = O\left(\frac{\sqrt{|\mathcal{C}| \log |\mathcal{C}|}}{m}\right).$$

$|\mathcal{C}|$ is at most n , so this bound dominates Theorem 3.1. In 10 of the 11 real-world instances we study, $|\mathcal{C}|$ is also smaller than k^2 (Appendix A), in which case this bound also dominates Theorem 3.2.

At a high level, our beyond-worst-case upper bounds are obtained not by rounding p , but instead using the structure of the sortition instance to abstract the problem into one about “types.” For this bound we then solve an LP in terms of “types,” round that LP, and then reconstruct a rounded panel distribution $\bar{p}, \bar{\pi}$ from the “type” solution. In particular, the *types* of individuals are the feature vectors which appear in the pool, and *types* of panels are the multisets of k feature vectors that satisfy the instance quotas. Fixing an instance, we project some p into type space by viewing it as a distribution \mathbf{p} over types of panels \mathfrak{K} , inducing marginals τ_c for each type individuals $c \in \mathcal{C}$.

In order to solve and round, we define the TYPE LP, which is analogous to Eq. (3.1). We let Q be the type analog of M , so that entry Q_{cj} is the number of individuals i with $F(i) = c$ contained in panels of type $j \in \mathfrak{K}$.⁵ Then,

$$\text{TYPE LP} \quad Q \mathbf{p} = \tau \tag{3.3}$$

$$\|\mathbf{p}\|_1 = 1 \tag{3.4}$$

$$\mathbf{p} \geq 0.$$

⁴The class of all anonymous marginals π includes the maximizers π^* of all reasonable fairness objectives, and second, this condition is satisfied by all existing selection algorithms used in practice, to our knowledge.

⁵Completing the analogy, $\mathcal{C}, \mathfrak{K}, Q, \mathbf{p}, \bar{p}, \tau$ are the “type” versions of $N, \mathcal{K}, M, p, \bar{p}, \pi$ from the original LP.

We round p in this LP to a panel type distribution \bar{p} while preserving (3.4). All that remains, then, is to construct some $\bar{p}, \bar{\pi}$ such that p is consistent with \bar{p} and $\|\pi - \bar{\pi}\|_\infty$ is small. This \bar{p} is in general supported by panels outside of $supp(p)$, unlike the \bar{p} obtained by Theorem 3.1. It is the anonymity of π which allows us to construct these new panels and prove that they are feasible for the instance.

3.3 Lower Bounds

This method of using bounded discrepancy to derive nearly fairness-optimal uniform lotteries has its limits, since there are even sparse M and fractional x for which no integer \bar{x} yields nearby $M\bar{x}$. In the worst case, we establish lower bounds by modifying those of Beck and Fiala [Spe85]:

Theorem 3.4. *There exist p, π for which for all uniform lotteries $\bar{p}, \bar{\pi}$,*

$$\min_{\bar{p} \in \mathcal{D}} \|\pi - \bar{\pi}\|_\infty = \Omega\left(\frac{\sqrt{k}}{m}\right).$$

Our k -dependent upper and lower bounds are separated by a factor of \sqrt{k} , matching the current upper and lower bounds of the Beck-Fiala conjecture as applied to linear discrepancy (also known as the lattice approximation problem [Sri97]). The respective gaps are incomparable, however, since for a given $x \in [0, 1]^n$, the former problem aims to minimize $\|M(x - \bar{x})\|_\infty$ over $\bar{x} \in \{0, 1\}^n$, while we aim to do the same over the $\bar{x} \in \mathbb{Z}^n$ for which $\sum_j x_j = \sum_j \bar{x}_j$ (Lemma B.2).

4 Theoretical Bounds on Fairness Loss

Since the fairness of a distribution p is determined by its marginals π , it is intuitive that if uniform lotteries incur only small marginal discrepancy (per Section 3), then they should also incur only small fairness losses. This should hold for any fairness notion that is sufficiently “smooth” (i.e., doesn’t change too quickly with changing marginals) in the vicinity of p, π .

Although our bounds from Section 3 apply to any reasonable initial distribution p , we are particularly concerned with bounding fairness loss from *maximally fair* initial distributions p^* . Here, we specifically consider such p^* that are optimal with respect to Maximin and NW. We note that, since there exist anonymous p^*, π^* that maximize these objectives, we can apply any upper bound from Section 3 to upper bound $\|\pi^* - \bar{\pi}\|_\infty$ (for any $\bar{\pi}$). We defer omitted proofs to Appendix C.

4.1 Maximin

Since Leximin is the objective optimized in practice, it would be most natural start with a p^* that is Leximin-optimal and bound loss with respect to this objective. However, the fact that Leximin fairness cannot be summarized by a single expression prevents us from formulating such a bound. Thus, we first pursue bounds on the closely-related objective, Maximin. We note that smallest marginal given by Leximin is in general not attained by any uniform lottery \bar{p} , and so in the most meaningful sense, a worst-case Maximin guarantee is a Leximin guarantee.

First, we show there exists some $\bar{p}, \bar{\pi}$ that gives bounded Maximin loss from p^*, π^* , the Maximin-optimal unconstrained distribution. This bound follows from Theorems 3.2 and B.6, using the simple observation that \bar{p} can decrease the lowest marginal given by p^* by no more than $\|\pi^* - \bar{\pi}\|_\infty$. Here $n_{min} := \min_c n_c$ denotes the smallest number of individuals which share any feature vector $c \in \mathcal{C}$.

Corollary 4.1. *By Theorem 3.2 and Theorem 3.3, for Maximin-optimal p^* , there exists a uniform lottery \bar{p} that satisfies*

$$\text{Maximin}(p^*) - \text{Maximin}(\bar{p}) = \frac{1}{m} \cdot O\left(\min\left\{\sqrt{|\mathcal{C}| \log |\mathcal{C}|}, \frac{k}{n_{min}} + 1\right\}\right).$$

Theorem 3.4 demonstrates that we cannot get an upper bound on Maximin loss stronger than $O(\sqrt{k}/m)$ using a uniform bound on changes in all π_i . However, since Maximin is concerned only with the smallest π_i , it seems plausible that better upper bounds on Maximin loss could result from rounding π while tightly controlling only losses in the smallest π_i ’s, while giving freer reign to larger marginals. We show that this is not the case by further modifying the instances from Theorem 3.4 to obtain the following lower bound on the Maximin loss:

Theorem 4.1. *There exists a Maximin-optimal p^* such that, for all uniform lotteries \bar{p} ,*

$$\text{Maximin}(p^*) - \text{Maximin}(\bar{p}) = \Omega\left(\frac{\sqrt{k}}{m}\right).$$

4.2 Nash Welfare

As NW has also garnered interest by practitioners and is applicable in practice [HG20], we upper-bound the NW fairness loss. Unlike Maximin loss, an upper bound on NW loss does not immediately follow from one on $\|\pi - \bar{\pi}\|_\infty$, because decreases in smaller marginals have larger negative impact on the NW. As a result, the upper bound on NW resulting from Section 3 is slightly weaker than that on Maximin:

Theorem 4.2. *For NW-optimal p^* , there exists a uniform lottery \bar{p} that satisfies*

$$\text{NW}(p^*) - \text{NW}(\bar{p}) = \frac{k}{m} \cdot O\left(\min\left\{\sqrt{|\mathcal{C}| \log |\mathcal{C}|}, \frac{k}{n_{\min}} + 1\right\}\right).$$

We give an overview of the proof of Theorem 4.2. To begin, fix a NW-optimizing panel distribution p^*, π^* . Before applying our upper bounds on marginal discrepancy from Section 3, we must contend with the fact that if this bounded loss is suffered by already-tiny marginals, the NW may decrease substantially or even go to 0. Thus, we first prove Lemmas 4.1 and 4.2, which together imply that no marginal in π^* is smaller than $1/n$.

Lemma 4.1. *For NW-optimal p^* over a support of panels $\text{supp}(p^*)$, there exists a constant $\lambda \in \mathbb{R}^+$ such that, for all $P \in \text{supp}(p^*)$, $\sum_{i \in P} 1/\pi_i^* = \lambda$.*

Lemma 4.2. *For Nash-Welfare optimal p^*, π^* , we have that $\pi_i^* \geq 1/n$ for all $i \in N$.*

Lemma 4.1 follows from the fact that the partial derivative of NW with respect to the probability it assigns a given panel must be the same as that with respect to any other panel at p^* (otherwise, mass in the distribution could be shifted to increase the NW). Lemma 4.2 then follows by the additional observation that $\mathbb{E}_{P \sim p^*} [\sum_{i \in P} 1/\pi_i^*] = n$.

Finally Lemma 4.3 follows from the fact that Lemma 4.2 limits the potential multiplicative, and therefore additive, impact on the NW of decreasing any marginal by $\|\pi - \bar{\pi}\|_\infty$:

Lemma 4.3. *For NW-optimal p^*, π^* , there exists a uniform lottery $\bar{p}, \bar{\pi}$ that satisfies $\text{NW}(p^*) - \text{NW}(\bar{p}) \leq k \|\pi^* - \bar{\pi}\|_\infty$.*

As the NW-optimal marginals π^* are anonymous, we can apply the upper bounds given by Theorem 3.3 and Theorem B.6 to show the existence of a $\bar{p}, \bar{\pi}$ satisfying the claim of the theorem.

5 Practical Algorithms for Computing Fair Uniform Lotteries

Algorithms. First, we implement versions of two existing rounding algorithms, which are implicit in our worst-case upper bounds.⁶ The first is Pipage rounding [GKPS06], or PIPAGE, a randomized rounding scheme satisfying negative association [DJR07]. The second is BECK-FIALA, the dependent rounding scheme used in the proof of Theorem 3.2. To benchmark these algorithms against the highest level of fairness they could possibly achieve, we use integer programming (IP) to compute the fairest possible uniform lotteries over $\text{supp}(p^*)$, the panels over which p^* randomizes.⁷ We define IP-MAXIMIN and IP-NW to find uniform lotteries over $\text{supp}(p^*)$ maximizing Maximin and NW, respectively. We remark that the performance of these IPs is still subject to our theoretical upper and lower bounds. We provide implementation details in Appendix D.1.

One question is whether we should prefer the IPs or the rounding algorithms for real-world applications. Although IP-MAXIMIN appears to find good solutions at practicable speeds, IP-NW

⁶We do not implement the algorithm implicit in Theorem 3.3 because our results already present sufficient alternatives for finding excellent uniform lotteries in practice.

⁷Note that these lotteries are not necessarily universally optimal, as they can randomize over only $\text{supp}(p^*)$; conceivably, one could find a fairer uniform lottery by also randomizing over panels not in $\text{supp}(p^*)$. However, PIPAGE and BECK-FIALA are also restricted in this way, and thus must be weakly dominated by the IP.

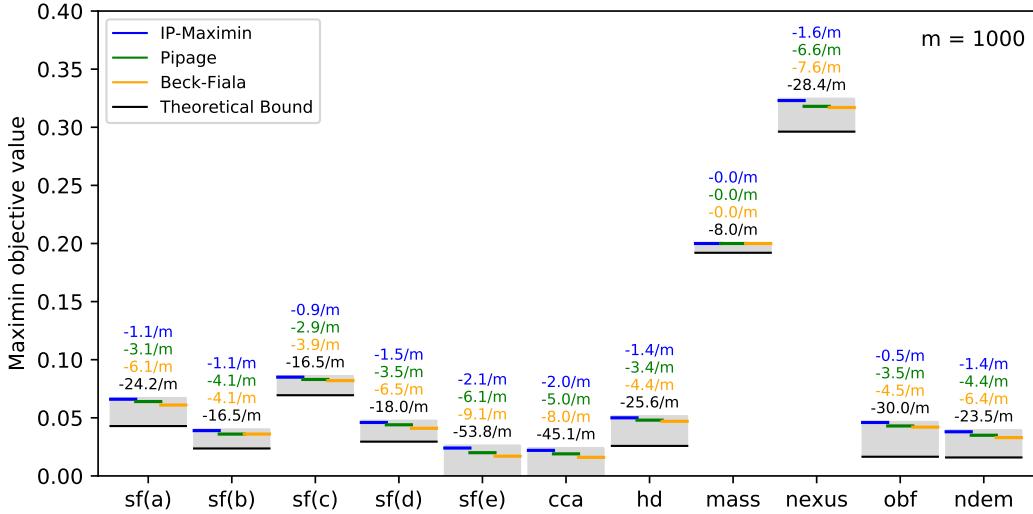


Figure 1: $m = 1000$. Shaded regions extend from $\text{Maximin}(p^*)$, the fairness of the optimal unconstrained distribution, down to the minimum fairness implied by the tightest theoretical upper bound in that instance (in all instances but ‘‘obf’’ Theorem 3.3 is tightest). Each algorithm or bound’s loss relative to $\text{Maximin}(p^*)$ is written above in the corresponding color. PIPAGE Maximin values are the average of 1000 runs (max standard deviation across instances was $< 10e-17$).

converges to optimality prohibitively slowly in some instances (see Appendix D.2 for runtimes). At the same time, we find that our simpler rounding algorithms give near-optimal uniform lotteries with respect to both fairness objectives. Also in favor of simpler rounding algorithms, many randomized rounding procedures (including Pipage rounding) have the advantage that they exactly preserve marginals over the combined steps of randomly rounding to a uniform lottery and then randomly sampling it—a guarantee that is much more challenging to achieve with IPs.

Uniform lotteries nearly exactly preserve Maximin, Nash Welfare fairness. We first measure the fairness of uniform lotteries produced by these algorithms in 11 real-world panel selection instances from 7 different organizations worldwide (instance details in Appendix A). In all experiments, we generate a lottery of size $m = 1000$. This is fairly small; it requires drawing only 3 balls from lottery machines, and in one instance we have that $m < n$. We nevertheless see excellent performance of all algorithms, and note that this performance will only improve with larger m .

Figure 1 shows the Maximin fairness of the uniform lottery computed by PIPAGE, BECK-FIALA, and IP-MAXIMIN for each instance. For intuition, recall that the level of Maximin fairness given by any lottery is exactly the minimum marginal assigned to any individual by that lottery. The upper edges of the gray boxes in Fig. 1 correspond to the optimal fairness attained by an unconstrained distribution p^* . These experiments reveal that the cost of transparency to Maximin-fairness is practically non-existent: across instances, the uniform lottery computed by IP-MAXIMIN decreases the minimum marginal by at most $2.1/m$, amounting to a loss of no more than 0.0021 in the minimum marginal probability in any instance. Visually, we can see that this loss is negligible relative to the original magnitude of even the smallest marginals given by p^* . Surprisingly, though PIPAGE and BECK-FIALA do not aim to optimize any fairness objective, they achieve only slightly larger losses in Maximin fairness, with PIPAGE outperforming BECK-FIALA. Finally, the heights of the gray boxes indicate that our theoretical bounds are often meaningful in practice, giving lower bounds on Maximin fairness well above zero in 9 out of 11 instances. We note these bounds only tighten with larger m . We present similarly encouraging results on NW loss in Appendix D.3.

Uniform lotteries nearly preserve all Leximin marginals. We still remain one step away from practice: our examination of Maximin does not address whether uniform lotteries can attain the finer-tuned fairness properties of the Leximin-optimal distributions currently used in practice. Fortunately, our results from Section 3 imply the existence of a uniform lottery that closely approximates

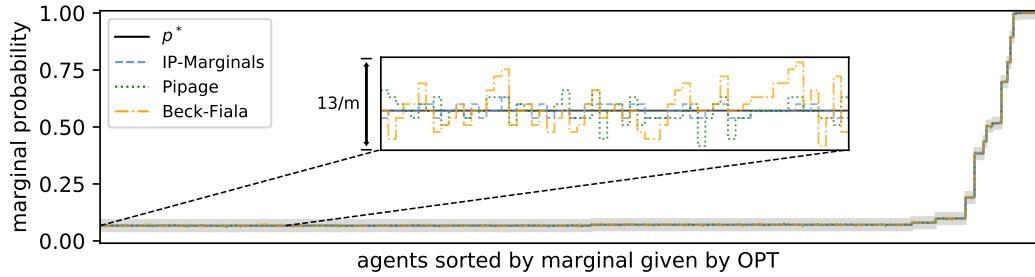


Figure 2: Instance = sf(a), $m = 1000$. Line plot shows the Leximin-optimal marginals π^* , along with marginals given by all algorithms sorted according to π^* . Note that each x coordinate then corresponds to an individual. The zoomed box shows the magnitude of marginal discrepancy around π^* . The surrounding shaded region shows the tightest theoretical bound on the marginal discrepancy, in this case from Theorem 3.3, around the optimal marginals. PIPAGE marginals are the average of 1000 runs (max standard deviation in any marginal was $< 10\text{e-}14$).

all marginals given by the Leximin-optimal distribution p^*, π^* . We evaluate the extent to which PIPAGE and BECK-FIALA preserve these marginals in Fig. 2. They are benchmarked against a new IP, IP-MARGINALS, which computes the uniform lottery over $\text{supp}(p^*)$ minimizing $\|\pi^* - \bar{\pi}\|_\infty$.

Figure 2 demonstrates that in the instance “sf(a)”, all algorithms produce marginals that deviate negligibly from those given by π^* . Analogous results on remaining instances appear in Appendix D.4 and show similar results. As was the case for Maximin, we see that our theoretical bounds are meaningful, but that we can consistently outperform them in real-world instances.

6 Discussion

Our aim was to show that uniform lotteries can preserve fairness, and our results ultimately suggest this, along with something stronger: that in practical instances, uniform lotteries can reliably almost exactly replicate *the entire set of marginals* given by the optimal unconstrained panel distribution. Our rounding algorithms can thus be plugged directly into the existing panel selection pipeline with essentially no impact on individuals’ selection probabilities, thus enabling translation of the output of Panelot (and other maximally fair algorithms) to a nearly maximally fair *and* transparent panel selection procedure. We note that our methods are not just compatible with ball-drawing lotteries, but any form of uniform physical randomness (e.g. dice, wheel-spinning, etc.).

Although we achieve our stated notion of transparency, a limitation of this notion is that it focuses on the final stage of the panel selection process. A more holistic notion of transparency might require that onlookers can verify that the panel is not being intentionally stacked with certain individuals. This work does not fully enable such verification: although onlookers can now observe individuals’ marginals, they still cannot verify that these marginals are *actually maximally fair* without verifying the underlying optimization algorithms. Thus, in the common case where quotas require even maximally fair algorithms to select certain individuals with probability near 1, onlookers cannot distinguish maximally fair selection from the scenario in which someone is being intentionally placed on the panel.

In research on economics, fair division, and other areas of AI, randomness is often proposed as a tool to make real-world systems fairer [Gri04, BCKM13, FSV20]. Nonetheless, in practice, these systems (with a few exceptions, such as school choice [New21]) remain stubbornly deterministic. Among the hurdles to bringing the theoretical benefits of randomness into practice is that allocation mechanisms fare best when they can be readily understood, and that randomness can be perceived as undesirable or suspect. Sortition is a rather unique paradigm at the heart of this tension: it relies centrally on randomness, while in the public sphere it is attaining increasing political influence. It is therefore a uniquely high-impact domain in which to study how to combine the benefits of randomness, such as fairness, with transparency. We hope that this work and its potential for impact will inspire the investigation of fairness-transparency tradeoffs in other AI applications.

References

- [BCKM13] Eric Budish, Yeon-Koo Che, Fuhito Kojima, and Paul Milgrom. Designing random allocation mechanisms: theory and applications. *American Economic Review*, 103(2):585–623, 2013.
- [BF81] József Beck and Tibor Fiala. “Integer-making” theorems. *Discrete Applied Mathematics*, 3(1):1–8, 1981.
- [BFT12] Dimitris Bertsimas, Vivek F. Farias, and Nikolaos Trichakis. On the efficiency-fairness trade-off. *Management Science*, 58(12):2234–2250, 2012.
- [BGP19] Gerdus Benadè, Paul Götz, and Ariel D. Procaccia. No stratification without representation. In *Proceedings of the 2019 ACM Conference on Economics and Computation (EC)*, pages 281–314, 2019.
- [BHJ⁺18] Richard Berk, Hoda Heidari, Shahin Jabbari, Michael Kearns, and Aaron Roth. Fairness in criminal justice risk assessments: The state of the art. *Sociological Methods & Research*, 50:3–44, 2018.
- [BJ12] Petter Brändén and Johan Jonasson. Negative dependence in sampling. *Scandinavian Journal of Statistics*, 39(4):830–838, 2012.
- [CELM07] Yann Chevaleyre, Ulle Endriss, Jérôme Lang, and Nicolas Maudet. A short introduction to computational social choice. In *Proceedings of the 33rd Conference on Current Trends in Theory and Practice of Computer Science (SOFSEM)*, pages 51–69, 2007.
- [DJR07] Devdatt Dubhashi, Johan Jonasson, and Desh Ranjan. Positive influence and negative dependence. *Combinatorics, Probability & Computing*, 16(1):29, 2007.
- [Doe07] Benjamin Doerr. Roundings respecting hard constraints. *Theory of Computing Systems*, 40(4):467–483, 2007.
- [DR96] Devdatt Dubhashi and Desh Ranjan. Balls and bins: A study in negative dependence. *BRICS Report Series*, 3(25):30, 1996.
- [FGG⁺20] Bailey Flanigan, Paul Götz, Anupam Gupta, Ariel D. Procaccia, and Gili Rusak. Panelot, 2020. <http://www.panelot.org/>.
- [FGG⁺21] B. Flanigan, P. Götz, A. Gupta, B. Hennig, and A. D. Procaccia. Fair algorithms for selecting citizens’ assemblies. *Nature*, 2021. Forthcoming.
- [FGGP20] Bailey Flanigan, Paul Götz, Anupam Gupta, and Ariel D Procaccia. Neutralizing self-selection bias in sampling for sortition. In *Proceedings of the 34th Annual Conference on Neural Information Processing Systems (NeurIPS)*, pages 6528–6539, 2020.
- [FKL16] Benjamin Fish, Jeremy Kun, and Ádám D Lelkes. A confidence-based approach for balancing fairness and accuracy. In *Proceedings of the 2016 SIAM International Conference on Data Mining (ICDM)*, pages 144–152, 2016.
- [FSV20] Rupert Freeman, Nisarg Shah, and Rohit Vaish. Best of both worlds: ex-ante and ex-post fairness in resource allocation. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*, pages 21–22, 2020.
- [GKPS06] Rajiv Gandhi, Samir Khuller, Srinivasan Parthasarathy, and Aravind Srinivasan. Dependent rounding and its applications to approximation algorithms. *Journal of the ACM (JACM)*, 53(3):324–360, 2006.
- [Gri04] Geoffrey Grimmett. Stochastic apportionment. *The American Mathematical Monthly*, 111(4):299–307, 2004.
- [HG20] Brett Hennig and Paul Götz. Sortition foundation stratification application, 2020. <https://github.com/sortitionfoundation/stratification-app>.

[ICAP19] The Irish Citizens’ Assembly Project, 2019. <http://www.citizenassembly.ie/work/>.

[Mou03] Hervé Moulin. *Fair division and collective welfare*. MIT press, 2003.

[MP16] Sam J. Maglio and Evan Polman. Revising probability estimates: Why increasing likelihood means increasing impact. *Journal of Personality and Social Psychology*, 111(2):141, 2016.

[New21] New York City Department of Education. Random selection in admissions, 2021. <https://www.schools.nyc.gov/enrollment/enroll-grade-by-grade/how-students-get-offers-to-doe-public-schools/random-numbers-in-admissions>.

[OEC20] OECD. *Innovative Citizen Participation and New Democratic Institutions*. OECD Publishing, 2020.

[RB08] Valerie F Reyna and Charles J Brainerd. Numeracy, ratio bias, and denominator neglect in judgments of risk and probability. *Learning and Individual Differences*, 18(1):89–107, 2008.

[Spe85] Joel Spencer. Six standard deviations suffice. *Transactions of the American Mathematical Society*, 289(2):679–706, 1985.

[Sri97] Aravind Srinivasan. Improving the discrepancy bound for sparse matrices: Better approximations for sparse lattice approximation problems. In *Proceedings of the 8th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 692–701, 1997.

[TBG⁺20] Andrew Trask, Emma Bluemke, Ben Garfinkel, Claudia Ghezzou Cuervas-Mons, and Allan Dafoe. Beyond privacy trade-offs with structured transparency. arXiv 2012.08347, 2020.

[WML20] Sean Jeremy Westwood, Solomon Messing, and Yphtach Lelkes. Projecting confidence: How the probabilistic horse race confuses and demobilizes the public. *The Journal of Politics*, 82(4):1530–1544, 2020.

A Panel Selection Datasets

We examine data from the following 11 real-world sortition panel selection instances, generously provided to us by several groups that specialize in organizing citizens’ assemblies. Table 1 shows the instance short-names we use throughout the paper, and which organization was responsible for each panel. The final two columns compare the values of our theoretical upper bounds on the marginal discrepancy, illustrating that in all instances except “obf”, the bound from Theorem 3.3 is tighter. Finally, we give some metadata about each instance, which is required for calculating the values of our theoretical upper bounds.

In particular, n = number of pool members, k = number panel members, \mathcal{C} = set of distinct realized feature-vectors in the pool. Precise constants used for computing exact the upper bounds are derived in Appendix B: the Theorem 3.2 bound is exactly k/m , the Theorem 3.3 bound is exactly

$$\frac{\sqrt{\frac{1}{2}(1 + \frac{\ln 2}{\ln |\mathcal{C}|})} \cdot \sqrt{|\mathcal{C}| \ln(|\mathcal{C}|)} + 1}{m},$$

and the Theorem B.6 bound is exactly $\frac{2k/n_{min}+1}{m}$. In all instances, $n_{min} = 1$.

B Omitted Proofs and Additional Beyond-Worst-Case Upper Bounds from Section 3

B.1 Omitted Proofs

We will make repeated use of the following generalization of Hoeffding’s inequality (see eg. Proposition 5 of [DR96]):

Table 1: Instance parameters and resulting theoretical bounds

Instance	Organization	n	k	$ \mathcal{C} $	Thm 3.2	Thm 3.3	Thm B.6
sf(a)	Sortition Foundation	312	35	182	$35/m$	$24.2/m$	$71/m$
sf(b)	Sortition Foundation	250	20	92	$20/m$	$16.5/m$	$41/m$
sf(c)	Sortition Foundation	161	44	92	$44/m$	$16.5/m$	$89/m$
sf(d)	Sortition Foundation	404	40	108	$40/m$	$18.0/m$	$81/m$
sf(e)	Sortition Foundation	1727	110	762	$110/m$	$53.8/m$	$221/m$
cca	Center for Climate Assemblies	825	75	554	$75/m$	$45.1/m$	$151/m$
hd	Healthy Democracy	239	30	202	$30/m$	$25.6/m$	$61/m$
mass	MASS LBP	70	24	25	$24/m$	$8.0/m$	$49/m$
nexus	Nexus	342	170	242	$170/m$	$28.4/m$	$341/m$
obf	Of By For	321	30	294	$30/m$	$31.6/m$	$61/m$
ndem	New Democracy	398	40	173	$40/m$	$23.5/m$	$81/m$

Lemma B.1. *If $\{\xi_j\}$ are negatively associated random variables with $\xi_j \in [a_j, b_j]$ and $\xi = \sum_j \xi_j$, then*

$$\Pr[|\mathbb{E}[\xi] - \xi| \geq t] \leq 2 \exp\left\{-\frac{2t^2}{\sum_j (b_j - a_j)^2}\right\}.$$

Here is our first use:

Theorem 3.1. *For any realizable π , we may efficiently randomly generate \bar{p} such that its marginals $\bar{\pi}$ satisfy*

$$\|\pi - \bar{\pi}\|_\infty = O\left(\frac{\sqrt{n \log n}}{m}\right).$$

Proof of Theorem 3.1. Given a vector of marginals π , let p be a basic solution to $Mp = \pi$, where M is the individual-feasible panel membership matrix, so that $|\text{supp}(p)| \leq n$.

To construct our desired uniform lottery \bar{p} , we will randomly sample a collection of m panels denoted by the vector $x \in \mathbb{Z}_+^{|\mathcal{K}|}$, where x_j is the number of copies of panel j in our set. We will do this sampling in such a way that $E[x_j] = m \cdot p_j$. Once x is constructed, we will let $\bar{p} := x/m$.

We first define the vectors $\lfloor mp \rfloor$ and $\tilde{m}p$, where $\lfloor mp \rfloor_j = \lfloor mp_j \rfloor$ and $\tilde{m}p = mp - \lfloor mp \rfloor$. First, to x we deterministically add $\lfloor mp \rfloor_j$ copies of each panel j . Then, there are $m - \|\lfloor mp \rfloor\|_1$ spots left, which we will fill by randomly sampling $\tilde{x} \in \{0, 1\}^{\mathcal{K}}$ such that $\mathbb{E}[\tilde{x}_j] = \tilde{m}p_j$ for all panels j . We additionally wish to sample these \tilde{x}_j such that $\|\tilde{x}\|_1 = m - \|\lfloor mp \rfloor\|_1$ and the \tilde{x}_j are *negatively associated*, as defined in [BJ12, DR96]. This may be accomplished via any number of sampling algorithms [BJ12].

We will next analyze the marginal $\bar{\pi}_i$ provided to any given individual i by this randomly sampled $\bar{p} = x/m$. Consider the collection of \tilde{x}_j for which i is contained in panel j . Note that these \tilde{x}_j are also negatively associated. Then for any $t \geq 0$,

$$\Pr[|\pi_i - \bar{\pi}_i| \geq t/m] = \Pr[|m\pi_i - m\bar{\pi}_i| \geq t] \tag{B.1}$$

$$= \Pr\left[\left|\mathbb{E}\left[\sum_{j \ni i} \tilde{x}_j\right] - \sum_{j \ni i} \tilde{x}_j\right| \geq t\right], \tag{B.2}$$

by the definition of x_j and $\mathbb{E}[\tilde{x}_j]$. Then by Hoeffding (Lemma B.1),

$$\leq 2 \exp\left(\frac{-2t^2}{|\{j : i \in j\}|}\right) \tag{B.3}$$

$$\leq 2 \exp\left(\frac{-2t^2}{n}\right), \tag{B.4}$$

where here we use that $|supp(p)| \leq n$. Then taking $t = \sqrt{\frac{1+\epsilon}{2} n \log n}$,

$$\leq \frac{2}{n^{1+\epsilon}}. \quad (\text{B.5})$$

Union bounding over all n individuals, we find that

$$\Pr \left[\|\pi - \bar{\pi}\|_\infty \geq \sqrt{(1+\epsilon)/2} \frac{\sqrt{n \log n}}{m} \right] \leq \frac{2}{n^\epsilon} < 1.$$

□

Note: if we are additionally guaranteed that all of the $\pi_i = \Omega(k/n)$, then a multiplicative form of Chernoff yields

$$\|\pi - \bar{\pi}\|_\infty = O \left(\sqrt{\frac{k \log n}{mn}} \right)$$

with constant probability.

Theorem 3.2. *For any realizable π , we may efficiently construct \bar{p} such that its marginals $\bar{\pi}$ satisfy*

$$\|\pi - \bar{\pi}\|_\infty \leq k/m.$$

Proof of Theorem 3.2. This result follows from the proof of Lemma 9 in Appendix B.4.1 of Flanigan *et al* [FGGP20]. Their proof invokes an algorithm derived from a notable theorem by Beck and Fiala [BF81], with a slight modification that ensures the exact preservation of their adding up constraint. In our setting, this modification ensures that \bar{p} , the rounded version of p , is also a valid distribution.

We note that rather than directly rounding p , we apply their argument to round a *modified version* of p , call it p' . We make this modification because p must be rounded to multiples of $1/m$, and their approach does 0/1 rounding. The vector p' is defined so it can be rounded to a 0/1 vector as follows: for all panels j , $p'_j = p_j \cdot m - \lfloor p_j \cdot m \rfloor$. Then, the algorithm rounds the entries of p' to a 0/1 vector \bar{p}' such that $\|p'\| = \|\bar{p}'\|$. With this rounding complete, the final uniform lottery \bar{p} is expressed as $\bar{p}_j = \bar{p}'_j / m + \lfloor p_j \cdot m \rfloor / m$. Observe that \bar{p}_j is equal an integer divided by m , and thus the resulting lottery \bar{p} contains multiples of $1/m$. One can also verify that the adding up constraint is conserved:

$$\sum_{j \in \mathcal{K}} \bar{p}_j = \sum_{j \in \mathcal{K}} (\bar{p}'_j / m + \lfloor p_j \cdot m \rfloor / m) = \sum_{j \in \mathcal{K}} (p'_j / m + \lfloor p_j \cdot m \rfloor / m) = \sum_{j \in \mathcal{K}} p_j$$

Finally, we note that the upper bound of $|F|$ on discrepancy that emerges from Flanigan *et al*'s argument corresponds to a discrepancy of k in our setting. Then, when $\bar{p}', \bar{\pi}'$ are re-scaled by $1/m$ to produce the final uniform lottery $\bar{p}, \bar{\pi}$, this bound translates to a maximum discrepancy of k/m in any marginal, giving our $\|\pi - \bar{\pi}\|_\infty \leq k/m$ bound.

□

Theorem 3.3. *If π is anonymous and realizable, then we may efficiently construct \bar{p} such that its marginals $\bar{\pi}$ satisfy*

$$\|\pi - \bar{\pi}\|_\infty = O \left(\frac{\sqrt{|\mathcal{C}| \log |\mathcal{C}|}}{m} \right).$$

Proof of Theorem 3.3. We begin with anonymous marginals π witnessed by some distribution p over \mathcal{K} . The first order of business is to project p into “type space,” in order to derive a distribution over panel types. Overloading F , we let $F(P) = \mathfrak{P}$ denote the panel type of a given panel P , defined as the multiset $F(P) = \{F(i) : i \in P\}$. Then we define the distribution over panel types induced by p as \mathfrak{p} , where the probability of drawing panel type \mathfrak{P} from \mathfrak{p} is defined as $\mathfrak{p}_{\mathfrak{P}} := \sum_{P \in \mathcal{K}: F(P) = \mathfrak{P}} p_P$.

This \mathfrak{p} satisfies the PANEL TYPE LP in Eq. (3.3). As an aside, note that this \mathfrak{p} has support $supp(\mathfrak{p}) = \{F(P) : P \in supp(p)\}$. We will assume without loss of generality that \mathfrak{p} is a basic solution to (3.3), so that it has at most $|\mathcal{C}|$ nonzero entries, where \mathcal{C} is the set of all feature-vectors appearing in the

pool, i.e., $supp(p) \leq |\mathcal{C}|$. Since $|supp(p)| \leq n$ without loss of generality, $|supp(\mathbf{p})| \leq n$ also, and so this basic \mathbf{p} may be found efficiently.

Given this distribution \mathbf{p} over panel types, we will round it to a uniform lottery $\bar{\mathbf{p}}$ of size m over panel types \mathcal{K} . Finally, we will lift this distribution over panel types $\bar{\mathbf{p}}$ back to a distribution \bar{p} over panels with the desired guarantee, and argue that this lift can be performed when the original marginals π are anonymous.

We randomly generate $\bar{\mathbf{p}}$ from \mathbf{p} via randomized rounding analogous to what was deployed in Theorem 3.1. We will again do this in such a way that $\mathbb{E}[x_j] = m \cdot \mathbf{p}_j$ for all panel types j . Once x is constructed, we will let $\bar{\mathbf{p}} := x/m$.

We again define the vectors $\lfloor m\mathbf{p} \rfloor$ and $\tilde{m}\mathbf{p}$, where $\lfloor m\mathbf{p} \rfloor_j = \lfloor m\mathbf{p}_j \rfloor$ and $\tilde{m}\mathbf{p} = m\mathbf{p} - \lfloor m\mathbf{p} \rfloor$. First, to x we deterministically add $\lfloor m\mathbf{p} \rfloor_j$ copies of each panel type j . Then, there are $m - \|\lfloor m\mathbf{p} \rfloor\|_1$ spots left, which we will fill by randomly sampling $\tilde{x} \in \{0, 1\}^{|\mathcal{K}|}$ such that $\mathbb{E}[\tilde{x}_j] = \tilde{m}\mathbf{p}_j$ for all panel types j . We additionally wish to sample these \tilde{x}_j such that $\|\tilde{x}\|_1 = m - \|\lfloor m\mathbf{p} \rfloor\|_1$ and the \tilde{x}_j are negatively associated, which is again doable.

Recall that type marginals $\tau_c, \bar{\tau}_c$ represent the expected number of panel spots allocated to each feature vector c by $\mathbf{p}, \bar{\mathbf{p}}$, respectively. We will next analyze the proximity of the rounded type marginals $\bar{\tau}_c$ to the original type marginals τ_c . Proceeding via an analysis similar to that of Theorem 3.1, we consider the collection of \tilde{x}_j for which feature vector c appears on panel type j (i.e., $Q_{cj} > 0$). We note that these \tilde{x}_j are again negatively associated, and thus all $Q_{cj}\tilde{x}_j$ are negatively associated, since for a fixed instance all Q_{cj} are constant.

Then for any $t \geq 0$,

$$\Pr[|\tau_c - \bar{\tau}_c| \geq t/m] = \Pr[|m\tau_c - m\bar{\tau}_c| \geq t] \quad (\text{B.6})$$

$$= \Pr\left[\left|\mathbb{E}\left[\sum_j Q_{cj}\tilde{x}_j\right] - \sum_j Q_{cj}\tilde{x}_j\right| \geq t\right], \quad (\text{B.7})$$

by the definition of x_j and \tilde{x}_j . Then by Hoeffding (Lemma B.1) with $\xi_j = Q_{cj}\tilde{x}_j$,

$$\leq 2 \exp\left(\frac{-2t^2}{\sum_j Q_{cj}^2}\right) \quad (\text{B.8})$$

$$\leq 2 \exp\left(\frac{-2t^2}{|\mathcal{C}|m_c^2}\right), \quad (\text{B.9})$$

where $m_c := \max_j Q_{cj}$, because $|supp(p)| \leq |\mathcal{C}|$. Therefore taking $t_c = \alpha \cdot m_c \cdot \sqrt{|\mathcal{C}| \log |\mathcal{C}|}$,

$$\leq \frac{2}{|\mathcal{C}|^{2\alpha^2}}. \quad (\text{B.10})$$

Taking $\alpha > \sqrt{\frac{1}{2}(1 + \frac{\log 2}{\log |\mathcal{C}|})}$ and union bounding over all $|\mathcal{C}|$ feature vectors, we may guarantee that with positive probability, the $\bar{\tau}$ resulting from this $\bar{\mathbf{p}}$ satisfies

$$\|\tau - \bar{\tau}\|_\infty \leq \alpha \cdot m_c \frac{\sqrt{|\mathcal{C}| \log |\mathcal{C}|}}{m}. \quad (\text{B.11})$$

Given such a $\bar{\tau}$, it remains to construct some uniform lottery \bar{p} over \mathcal{K} which is consistent with $\bar{\tau}$ and additionally satisfies the desired guarantees on $\bar{\tau}$. The challenges here are to ensure that

1. $0 \leq \bar{\pi}_i \leq 1$ for all i ,
2. each individual appears on each panel in \bar{p} at most once,⁸ and
3. $|\pi_i - \bar{\pi}_i|$ is small for all i .

⁸We note that this is a concern because we will not simply be choosing panels from collection \mathcal{K} (as we don't see the entire collection *a priori*); we will instead be *constructing* panels that must turn out to be feasible.

We will describe an algorithm for forming \bar{p} from \bar{p} and argue that it satisfies all three of these criteria.

In particular, \bar{p} encodes a lottery over m panel types \mathfrak{P}_j , possibly with multiplicity. From each panel type \mathfrak{P}_j in the support of \bar{p} , we form one panel in the support of \bar{p} by, for each $c \in \mathcal{C}$, allocating each of panel type \mathfrak{P}_j 's Q_{cj} "spots" to individuals $i \in N_c$.

We will populate each panel type \mathfrak{P}_j with individuals for each c independently. Once this is done for all $c \in \mathcal{C}$, we will have a panel $P_j \in \text{supp}(\bar{p})$. Therefore all of what follows is for fixed $c \in \mathcal{C}$, and will be done for each c in turn. For convenience, let $N_c := \{i \in [n] : F(i) = c\}$ be the n_c individuals with feature vector c .

For each $i \in N_c$, we would like to take $\bar{\pi}_i = \bar{\tau}_c/n_c$. (This is the aim, since π is anonymous by assumption, and so $\pi_i = \tau_c/n_c$.)

Since $\bar{\tau}_c/n_c$ may not be an integer, round up or down by 1 such that $\sum_{i \in N_c} \bar{\pi}_i = \bar{\tau}_c$ arbitrarily. Clearly these $\bar{\pi}_i \geq 0$. To see that $\bar{\pi}_i \leq 1$, observe that since \mathfrak{p} is a distribution,

$$\bar{\tau}_c = \sum_j \bar{p}_j Q_{cj} \leq \max_j Q_{cj} = m_c \leq n_c,$$

where the last inequality follows because all \mathfrak{P} are feasible panel types, so they cannot contain more individuals $i \in N_c$ than exist in the pool. Dividing by n_c yields $\bar{\pi}_i \leq \lceil \bar{\tau}_c/n_c \rceil \leq 1$. This satisfies (1).

Next we argue that these individuals $i \in N_c$ may be assigned to panels such that these marginals $\bar{\pi}_i$ are satisfied and no individual appears on any panel more than once. To do this, assign individuals to panels greedily, in the following sense. Order the m panel types \mathfrak{P}_j arbitrarily, and let

$$d_i^j := m \cdot \bar{\pi}_i - \sum_{j' < j} \mathbb{1}\{i \in P_{j'}\}$$

be the number of "spots" in \bar{p} of type c on which i still needs to be placed at the beginning of round j in order to reach their allocation of $m\bar{\pi}_i$ spots. (This d_i^j can be viewed as the "unsatisfied demand" of individual i at round j , according to the promised number of spots $m\bar{\pi}_i$.) At every step j , choose for panel type \mathfrak{P}_j the Q_{cj} individuals $i \in N_c$ with the largest d_i^j , ie. the most unsatisfied demand, to be on panel P_j .

Because the $\bar{\pi}_i$ are all either $\lfloor \bar{\tau}_c/n_c \rfloor$ or $\lceil \bar{\tau}_c/n_c \rceil$, the initial values of d_i^0 for $i \in N_c$ are all within 1 of one another. Note that the greedy algorithm preserves this property that d_i^j remain within 1 of one another for all steps j , since at each step j it decreases some collection of maximal d_i^j by 1. To see that the greedy algorithm succeeds, suppose for the sake of contradiction that we reach some first step j for which a position on panel P_j cannot be filled; then there are not enough individuals with remaining unmet demand, so $Q_{cj} > |\{i : d_i^j > 0\}|$. Since $Q_{cj} \leq m_c \leq n_c$, it must be the case that some $i \in N_c$ have already been fully assigned by step j (meaning that for them $d_i^j = 0$), and so all $d_i^j \in \{0, 1\}$ because the d_i^j are within 1 of one another. But $\sum_j Q_{cj} = \sum_i d_i^0 = m \cdot \bar{\tau}_c$, while at this point

$$\sum_{j' \geq j} Q_{cj'} \geq Q_{cj} > |\{i : d_i^j > 0\}| = \sum_i d_i^j,$$

meaning that the number of unallocated spots of type c remaining at step j exceeds the remaining unmet demand of the i . This implies that strictly more than Q_{cj} individuals i were given spots on panel j' for some earlier $j' < j$. But this cannot have happened, and so this greedy algorithm must succeed in feasibly assigning individuals of each type c to panels. Running it for each $c \in \mathcal{C}$ in turn gives a collection of quota-feasible panels, with no individual appearing on any panel more than once. This satisfies (2).

To conclude, we confirm that the individual marginals are close. By the anonymity of π , for all i with $F(i) = c$ we have $\pi_i = \tau_c/n_c$, and also $\bar{\pi}_i = \bar{\tau}_c/n_c \pm 1/m$. Since $m_c \leq n_c$, therefore (B.11) implies

$$|\pi_i - \bar{\pi}_i| \leq \frac{m_c}{n_c} \cdot \alpha \cdot \frac{\sqrt{|\mathcal{C}| \log |\mathcal{C}|}}{m} + \frac{1}{m} = O\left(\frac{\sqrt{|\mathcal{C}| \log |\mathcal{C}|}}{m}\right),$$

for all i , satisfying (3). \square

Theorem 3.4. *There exist p, π for which for all uniform lotteries $\bar{p}, \bar{\pi}$,*

$$\min_{\bar{p} \in \bar{\mathcal{D}}} \|\pi - \bar{\pi}\|_\infty = \Omega\left(\frac{\sqrt{k}}{m}\right).$$

Proof of Theorem 3.4. We will make use of the following lemma:

Lemma B.2. *Any k -uniform hypergraph on $[n]$ is realizable via quotas as the set of feasible panels for an instance of the panel selection problem with pool $[n]$.*

When individual membership in feasible panels is represented as $M \in \{0, 1\}^{n \times |\mathcal{K}|}$, this lemma claims that any M with uniform column norms is realizable by an instance of the panel selection problem.

Proof. Given a set system $\mathcal{S} \subseteq \binom{[n]}{k}$, we may construct a set of upper quotas such that the collection of feasible panels is exactly \mathcal{S} .

To do this, construct a binary feature f_T for each $T \notin \mathcal{S}$. For each i in $[n]$, let $f_T(i) = 1$ if and only if $i \in T$; otherwise let $f_T(i) = 0$. Finally, enforce the upper quota that for all feasible panels $P \subset [n]$,

$$\sum_{i \in P} f_T(i) \leq k - 1,$$

for all $T \notin \mathcal{S}$ —that is, no feasible panel has more than $k - 1$ members belonging to any T . Clearly no $T \notin \mathcal{S}$ is a feasible panel. For $S \in \mathcal{S}$, observe that $|S| = k$, and so for all $T \notin \mathcal{S}$, we have $|S \cap T| \leq k - 1$. Therefore all $S \in \mathcal{S}$ are feasible.

Finally, it bears noting that this is also possible to execute using lower quotas: taking $f'_T(i) = 1 - f_T(i)$, we could instead enforce for each $T \notin \mathcal{S}$ that

$$\sum_{i \in P} f'_T(i) \geq 1.$$

□

Using this lemma, our aim is to deploy some matrix $M \in \{0, 1\}^{n \times |\mathcal{K}|}$ for which

$$\min_{\bar{x} \in \bar{\Delta}} \|M\bar{x}\|_\infty = \Omega\left(\sqrt{k}\right),$$

where $\bar{\Delta} := \{x \in \{\dots, -3, -1, 1, 3, \dots\}^N : \sum_i x_i = 0\}$ and all columns of M sum to k . Translating and scaling appropriately and applying Lemma B.2, this will provide our desired $\Omega\left(\frac{\sqrt{k}}{m}\right)$ lower bound.

The common instances which provide lower bounds of $\Omega(\sqrt{k})$ for the Beck-Fiala problem are insufficient for our purposes in two respects. First, while they are column-sparse, they are generally not uniform in column norm. Second, they are incomparable in terms of the \bar{x} which they quantify over. The Beck-Fiala problem considers minimizing $\|M\bar{x}\|_\infty$ in the more restrictive rounding setting where $\bar{x} \in \{-1, 1\}^n$, while we are concerned with $\bar{x} \in \Delta$.

We overcome these barriers by first modifying the Walsh matrices — a family of Hadamard matrices — in order to guarantee uniform column norms, and then modifying the Beck-Fiala lower bound proof of [Spe85, Theorem 19] for arbitrary Hadamard matrices to apply to our matrices for all $\bar{x} \in (2\mathbb{Z} + 1)^n$.

To begin, let H_t be the $2^t \times 2^t$ Walsh matrix, defined recursively by $H_0 = 1$ and

$$H_{t+1} = \begin{bmatrix} H_t & H_t \\ H_t & -H_t \end{bmatrix}.$$

Let $N := 2^t$ denote its dimension. It is a fact that all rows (and columns) besides the first have an equal number of 1 and -1 entries. Therefore we take H'_t to be the submatrix derived by dropping the first two columns of H_t . (We remove the first column so that all remaining columns have equal

sum; we remove the second so that $\bar{\Delta}$ is nonempty). Additionally, let h_i denote the rows of H'_t , and h^j denote its columns. Then H'_t has the property that $\sum_i h_i^j = 0$, and in particular all columns h^j have $N/2$ 1-entries.

We have the following lemma:

Lemma B.3.

$$\min_{x \in \bar{\Delta}} \|H'_t x\|_\infty \geq \frac{N-2}{\sqrt{N}},$$

where $\bar{\Delta} := \{x \in \{\dots, -3, -1, 1, 3, \dots\}^{N-2}\}$.

Proof. This right-hand side is $H'_t x = (h_1 x, \dots, h_N x)^T$. We aim to show that there is some i for which $|h_i x|$ is large. Writing $\|H'_t x\|_2^2$ two ways, we have that

$$\begin{aligned} \sum_i (h_i x)^2 &= \|x_1 h^1 + \dots + x_{N-2} h^{N-2}\|_2^2 \\ &= \sum_j x_j^2 \|h^j\|_2^2 + \sum_{j \neq k} x_j x_k (h^j \cdot h^k). \end{aligned}$$

The entries of H_t are all ± 1 , and $h^j \cdot h^k = 0$ for $j \neq k$ (since the columns of H_t and therefore H'_t are orthogonal), so this becomes

$$\begin{aligned} &= (N-2) \sum_j x_j^2 \\ &\geq (N-2)^2, \end{aligned}$$

since $x_i^2 \geq 1$ by assumption. Therefore by averaging there is some i for which $(h_i x)^2 \geq \frac{(N-2)^2}{N}$, and so $|h_i x| \geq \frac{N-2}{\sqrt{N}}$, as desired. \square

Next we translate H'_t into an instance of the panel selection problem and argue it has the desired properties. Take $M := \frac{1}{2}(H_t + 1^{N \times (N-2)})$ to be the $\{0, 1\}$ matrix derived from H'_t .

The fact that M has uniform column norm $k = N/2$ directly follows from a property of Walsh matrices. Therefore we may apply Lemma B.2 to argue that M is realizable as the individual-panel membership matrix for some instance of the panel selection problem, with $n = N$, $|\mathcal{K}| = N-2$, and $k = N/2$.

To conclude, consider the uniform $p = \left(\frac{1}{N-2}, \dots, \frac{1}{N-2}\right)$, with $m = a(N-2) + (N-2)/2$ for any $a \in \mathbb{Z}_+$. In this case, each coordinate of p falls evenly between multiples of $1/m$ and must be rounded to multiples of $1/m$. Letting $x := p - \lfloor mp \rfloor / m = (1/2m, \dots, 1/2m)$ be this vector of remainders, we must replace it with some $\bar{x} \in (\mathbb{Z}/m)^{N-2}$, while maintaining that $\sum_j \bar{x}_j = \sum_j x_j = (N-2)/2m$, so that the resulting $\bar{p} = \lfloor mp \rfloor / m + \bar{x}$ remains a distribution over panels. (Note that here negative \bar{x}_j signify that the distribution mass on panel j decreases from p to \bar{p} .)

Explicitly, we then have

$$\|\pi - \bar{\pi}\|_\infty = \|Mp - M\bar{p}\|_\infty \tag{B.12}$$

$$= \|M(x - \bar{x})\|_\infty \tag{B.13}$$

$$= \frac{1}{2m} \|My\|_\infty, \tag{B.14}$$

where $y := 2m(\bar{x} - x)$.

$$= \frac{1}{2m} \left\| \frac{1}{2} H'_t y + \frac{1}{2} 1^{N \times (N-2)} y \right\|_\infty \tag{B.15}$$

$$= \frac{1}{4m} \|H'_t y\|_\infty, \tag{B.16}$$

where $\sum_i y_i = 0$ because we require that \bar{p} remain a distribution. Then since $y \in (2\mathbb{Z} + 1)^{N-2}$, by Lemma B.3 we have

$$\geq \frac{N-2}{4m\sqrt{N}} \tag{B.17}$$

$$= \Omega\left(\frac{\sqrt{k}}{m}\right), \tag{B.18}$$

since $k = N/2$.

This holds for all $y \in (2\mathbb{Z} + 1)^{N-2}$. Recall that $\bar{\mathcal{D}} := \{\bar{p} \in (\mathbb{Z}_+/m)^{|\mathcal{K}|} : \|\bar{p}\|_1 = 1\}$, and so

$$\bar{\mathcal{D}} \subseteq \{p + \bar{\Delta}/2m\}.$$

Therefore (B.18) implies that

$$\min_{\bar{p} \in \bar{\mathcal{D}}} \|\pi - \bar{p}\|_\infty = \Omega\left(\frac{\sqrt{k}}{m}\right),$$

as desired. \square

B.2 Additional Beyond-Worst-Case Upper Bounds

Since some of our beyond-worst-case upper bounds apply to anonymous realizable π , it is reasonable to ask how prevalent anonymous realizable π are, for arbitrary instances of sortition. Fortunately, we have the following claim:

Claim B.4. *For any instance of the panel selection problem and any realizable π , let π' be the “anonymized” marginals obtained by setting π'_i to the average $\pi_{i'}$ across all i' with the same feature vector as i . Then π' is realizable also.*

Proof of Claim B.4. Let π^* denote the “anonymization” of π , and take

$$\Pi := \left\{ \pi' : \text{realizable, and for all } c, \sum_{i:F(i)=c} \pi'_i = \sum_{i:F(i)=c} \pi_i \right\}.$$

We will show that $\pi^* \in \Pi$.

We argue by way of contradiction. Let $\hat{\pi}$ denote the “most anonymized” $\pi' \in \Pi$, in the sense that

$$\hat{\pi} = \arg \min_{\pi' \in \Pi} \max_c \left(\max_{i:F(i)=c} \pi'_i - \min_{i:F(i)=c} \pi'_i \right).$$

Let i and i' be some pair of individuals with $F(i) = F(i')$ witnessing this maximum diameter, and let p be a distribution with marginals $\hat{\pi}$. For each such pair, we will argue that p may be modified so that $\hat{\pi}_i = \hat{\pi}_{i'}$ while leaving all other marginals unchanged. By iteratively applying this to all such pairs, we will contradict the minimality of $\hat{\pi}$.

To start, observe that by assumption $\hat{\pi}_i > \hat{\pi}_{i'}$. Let p' be the distribution over feasible panels which is the same as p , except that i and i' switch places in any panel on which either of them appear. All such panel replacements yield feasible panels, since they have the same feature vector c . Finally take $p_{\text{new}} = (p + p')/2$. As promised, this distribution has the property that $\pi_i = \pi_{i'}$ and all other marginals are unchanged. \square

As a warm-up to the beyond-worst-case guarantees, we begin with the case when there is only one feature of interest, so that $F = \{f\}$. We place no constraints on the size of the set of feature values Ω , nor do we require that π is anonymous.

Theorem B.5. *If π is realizable and $|F| = 1$, then we may efficiently identify \bar{p} such that its marginals $\bar{\pi}$ satisfy*

$$\|\pi - \bar{\pi}\|_\infty < \frac{2}{m}.$$

Proof of Theorem B.5. Given marginals π , let p be a distribution over feasible panels \mathcal{K} which witnesses π . The first step of this rounding is to consider the marginals τ_v of each feature value v : $\tau_v = \sum_{i:f(i)=v} \pi_i$. Note that $\sum_v \tau_v = \sum_i \pi_i = k$. Since there is only one feature, all quotas for feasible panels P are of the form

$$l_v \leq |\{i \in P : f(i) = v\}| \leq u_v, \quad (\text{B.19})$$

and taking the expectation of this over p gives

$$l_v \leq \mathbb{E}_p[|\{i \in P : f(i) = v\}|] \leq u_v \quad (\text{B.20})$$

$$l_v \leq \tau_v \leq u_v. \quad (\text{B.21})$$

Therefore $l_v \leq \lfloor \tau_v \rfloor$ and $u_v \geq \lceil \tau_v \rceil$. We will construct a new distribution \bar{p} over panels P which satisfy $\lfloor \tau_v \rfloor \leq |\{i \in P : f(i) = v\}| \leq \lceil \tau_v \rceil$ for all features v , and are therefore guaranteed to be feasible.

We will construct feasible panels via the following scheme. Consider the interval $[0, km] \subset \mathbb{R}$ as representing the km spots to be allocated across the m panels which will comprise our lottery, and let $s_t := [t-1, t)$ denote spot t . Next observe that $m \sum_i \pi_i = km$, and so $m\pi_i$ may be viewed as the expected number of spots which p would give to i .

First group the π_i by feature value to form $\tau_v = \sum_{i:f(i)=v} \pi_i$, and then pack them into $[0, km]$, so that individuals with common feature values have contiguous sections; let S_i denote the portion of $[0, km]$ allocated to i , so that $|S_i| = \pi_i$. We will choose an individual $I(t)$ for each spot s_t , and then assemble the m panels that comprise \bar{p} by taking

$$P_r := \{I(t) : t = wm + r \text{ for } w \in \{0, \dots, k-1\}\}, \quad (\text{B.22})$$

for $r \in \{1, \dots, m\}$.

How to choose which individual will get the spot t for each t ? If $S_i \supseteq s_t$ then $I(t) = i$. Otherwise, s_t is split between two or more individuals, possibly with different feature values, in which case we call it *contested*. Observe that no matter how these contested s_t are allocated (no matter the choice of $I(t)$ for split t), it will be the case that $|\pi_i - \bar{\pi}_i| < 2/m$, since there is at most one contested s_t at each endpoint of the interval S_i .

It remains to argue that the panels chosen in (B.22) are feasible; in particular that $\lfloor \tau_v \rfloor \leq \bar{\tau}_v \leq \lceil \tau_v \rceil$ for all v . By construction, each panel P_r has some number of spots which will necessarily be allocated to an individual with feature value v , and some number of spots which are contested and may or may not be allocated to an individual with feature vector v . For each value v , there are at most two spots in all of $[0, km]$ which are *type-contested* in this way. If some panel P_r contains at most one type-contested spot for type v , then no matter which way it is allocated, $|\{i \in P : f(i) = v\}| - \tau_v| < 1$, and so P_r is feasible with respect to v . In the worst case, for some given v both of the spots which are type-contested by v appear on the same panel P_r . In order to ensure that $|\{i \in P : f(i) = v\}| - \tau_v| < 1$, it must be the case that exactly one of these two spots is allocated to some i for which $f(i) = v$. Fortunately this constraint is easily satisfiable, even in the case when a given panel P_r contains both of the type-contested spots for multiple features v .

Therefore the \bar{p} as constructed by (B.22) is supported by panels which are not only feasible but respect quotas which are maximally tight, given that the input p, π was realizable. Finally since each i contests at most two spots, we have that

$$\|\pi - \bar{\pi}\|_\infty < \frac{2}{m}.$$

□

Theorem B.6. *Given realizable anonymous π , we may efficiently identify $\bar{p}, \bar{\pi}$ such that*

$$\|\pi - \bar{\pi}\|_\infty = O\left(\frac{1}{m} \max\left\{\frac{k}{n_{\min}}, 1\right\}\right),$$

where $n_{\min} := \min_c n_c$ is the minimum number of individuals in the pool which share any one feature vector.

Proof. We proceed as in the proof of Theorem 3.3, but apply a different rounding to the panel type LP to obtain \bar{p} . To begin, p, π projects to some \mathbf{p}, τ . Without loss of generality assume that it is a basic solution to the TYPE LP, (3.4).

Again let $x \in \mathbb{Z}_+^{|\mathcal{R}|}$ be the vector of panel type counts which we select to form \bar{p} , so that $\bar{p} = x/m$ and $\|x\|_1 = m$. We will again start by deterministically choosing $\lfloor x \rfloor = \lfloor m\mathbf{p} \rfloor$, the integer part of $m\mathbf{p}$. In order to fill out $x = \lfloor x \rfloor + \tilde{x}$, we will once again construct $\tilde{x} \in \{0, 1\}^{|\mathcal{R}|}$ such that $\|\lfloor x \rfloor\|_1 + \|\tilde{x}\|_1 = m$. Let $\tilde{m}\mathbf{p} = m\mathbf{p} - \lfloor m\mathbf{p} \rfloor \in (0, 1)^{|\mathcal{R}|}$ denote the fractional solution which we round to obtain \tilde{x} .

Note that the constraint matrix Q in (3.3) has the property that for all columns q^j , $\|q^j\|_1 = k$. As a special case of [Doe07, Theorem 6], applied to $\tilde{m}\mathbf{p}$ and the panel type LP, there exists an $\tilde{x} \in \{0, 1\}^{|\mathcal{R}|}$ such that

$$\|Q(\tilde{m}\mathbf{p} - \tilde{x})\|_\infty < 2k.$$

and for which $\|\tilde{m}\mathbf{p}\|_1 = \|\tilde{x}\|_1$. (This follows from a generalization of the Beck-Fiala algorithm which both respects hard constraints and applies to arbitrary matrices Q with bounded column norms, and is therefore also algorithmic.)

It is clear that the resulting \bar{p} is a distribution over panel types. Its marginals therefore satisfy

$$\|\tau - \bar{\tau}\|_\infty = \|Q\mathbf{p} - Q\bar{p}\|_\infty = \frac{1}{m} \|Q(\tilde{m}\mathbf{p} - \tilde{x})\|_\infty < \frac{2k}{m}.$$

Given that such a $\bar{p}, \bar{\tau}$ exists, it remains to generate \bar{p} and $\bar{\pi}$ in such a way as to give the desired bound on the discrepancy in individual marginals. We proceed in a manner identical to the proof of Theorem 3.3.

Again we have that $\bar{\tau} \geq 0$ and $\bar{\tau} = \sum_j Q_{cj} \bar{p}_j \leq m_c \leq n_c$, where $m_c = \max_j Q_{cj}$ and n_c is the number of individuals i for which $F(i) = c$, since \bar{p} is a distribution over feasible panel types j . Therefore dividing $\bar{\tau}$ amongst the $\bar{\pi}_i$ as equally as possible for each c gives $\bar{\pi}_i \in [0, 1]$.

By the anonymity of π , for all i with $F(i) = c$, $\pi_i = \tau_c/n_c$, and dividing the spots in \bar{p} for feature vector c as equally as possible amongst the n_c individuals gives $\bar{\pi}_i \in \{\bar{\tau}_c/n_c \pm \frac{1}{m}\}$. This equal division of spots in order to form \bar{p} from $\bar{\pi}$ is feasible by the same greedy algorithm as in the proof of Theorem 3.3. Therefore the resulting $\bar{p}, \bar{\pi}$ satisfies

$$\begin{aligned} \|\pi - \bar{\pi}\|_\infty &= \max_c |\tau_c/n_c - \bar{\pi}_i| \\ &< \frac{1}{n_c} \|\tau - \bar{\tau}\|_\infty + \frac{1}{m} \\ &< \frac{2k}{n_{\min} \cdot m} + \frac{1}{m}. \end{aligned}$$

□

C Omitted Proofs from Section 4

Theorem 4.1. *There exists a Maximin-optimal p^* such that, for all uniform lotteries \bar{p} ,*

$$\text{Maximin}(p^*) - \text{Maximin}(\bar{p}) = \Omega\left(\frac{\sqrt{k}}{m}\right).$$

Proof of Theorem 4.1. We will follow the proof of Theorem 3.4: first we use the Walsh matrices to construct a matrix with the desired properties, prove a modified version of Lemma B.3 for it, and then appeal to Lemma B.2 to argue that it corresponds to a realizable instance of the panel selection problem.

In contrast to the construction in Theorem 3.4, where we need only demonstrate that *some* $\bar{\pi}_i$ deviates from π_i , we must construct an instance for which (essentially) the minimum π_i necessarily *decreases*. We accomplish this by first modifying the Walsh matrices to have uniform row norm, so that π is uniform and all π_i are minimal. We then introduce a second set of “twin” individuals, each

i' of which is a member of the panels which their twin i is not. This ensures that any discrepancy in $\bar{\pi} - \pi$ is witnessed in the downward direction.

To begin, again let H_t be the $2^t \times 2^t$ Walsh matrix, with $N := 2^t$ its dimension. This time we take H_t^* to be the submatrix derived by dropping the first row of H_t . By properties of Walsh matrices, all remaining rows in H_t^* have an equal number of 1 and -1 entries, (though this is no longer true of the columns).

Again letting h_i denote the rows of H_t^* , and h^j denote its columns, we have the following new version of Lemma B.3, which requires the additional assumption that $\sum_j x_j = 0$:

Lemma C.1.

$$\min_{x \in \Delta^*} \|H_t^* x\|_\infty \geq \sqrt{N},$$

where $\Delta^* := \{x \in \{\dots, -3, -1, 1, 3, \dots\}^N : \sum_j x_j = 0\}$.

Proof. This right-hand side is $H_t' x = (h_1 x, \dots, h_{N-1} x)^T$. We aim to show that there is some i for which $|h_i x|$ is large. Writing $\|H_t' x\|_2^2$ two ways, we have that

$$\begin{aligned} \sum_i (h_i x)^2 &= \|x_1 h^1 + \dots + x_N h^N\|_2^2 \\ &= \sum_j x_j^2 \|h^j\|_2^2 + \sum_{j \neq k} x_j x_k (h^j \cdot h^k) \end{aligned}$$

the entries of H_t^* are all ± 1 , and $h^j \cdot h^k = -1$ for $j \neq k$ (since the columns of H_t were orthogonal), so this becomes

$$\begin{aligned} &= (N-1) \sum_j x_j^2 - \sum_{j \neq k} x_j x_k \\ &= N \sum_j x_j^2 - \sum_j \sum_{k \neq j} x_j x_k \\ &= N \sum_j x_j^2 \\ &\geq N^2, \end{aligned}$$

since $x_i^2 \geq 1$ by assumption. Therefore by averaging there is some i for which $(h_i x)^2 \geq \frac{N^2}{N-1}$, and so $|h_i x| \geq \sqrt{N}$, as desired. \square

As constructed, all rows of H_t^* have the same number of 1s, so when we transform it into some M for some instance of the panel selection problem, it will yield that the marginals π of uniform p are uniform. However we cannot yet apply Lemma B.2, since the columns of the resulting M do not have constant norm; in particular, the first column will be all 1s.

In order to simultaneously correct for this and translate from ℓ_∞ to Maximin lower bounds, we introduce “twins” for each i . Letting $M^* = \frac{1}{2}(H_t^* + 1^{(N-1) \times N})$ be this $\{0, 1\}$ matrix, define $\bar{M}^* := 1^{(N-1) \times N} - M^*$ to be its complement, so that $M_{ij}^* = 1 - \bar{M}_{ij}^*$ for all i, j . Finally take

$$M = \begin{bmatrix} M^* \\ \bar{M}^* \end{bmatrix}$$

and observe that this $M \in \{0, 1\}^{(2N-2) \times N}$ has uniform column norm $N-1$ because of \bar{M}^* . We may therefore apply Lemma B.2 to claim that it is the individual-panel membership matrix of some instance of the panel selection problem.

The remainder of the argument proceeds similarly to that of Lemma B.3, with additional step of showing that the lower bound holds for the maximin objective. We include the full argument for completeness.

Similarly take $p = (\frac{1}{N}, \dots, \frac{1}{N})^T$, with $m = aN + N/2$ for any $a \in \mathbb{Z}_+$, $n = 2N-2$ (the number of individuals), and $k = N-1$. This p gives equal marginals: here $\pi_i = (Mp)_i = \frac{N-1}{2N-2} = \frac{k}{n}$

for all i . Again each coordinate of p falls evenly between multiples of $1/m$ and must be rounded to multiples of $1/m$. Letting $x := p - \lfloor mp \rfloor / m = (1/2m, \dots, 1/2m)^T$ be this vector of remainders, we must replace it with some $\bar{x} \in (\mathbb{Z}/m)^N$, while maintaining that $\sum_j \bar{x}_j = \sum_j x_j = N/2m$, so that the resulting $\bar{p} = \lfloor mp \rfloor / m + \bar{x}$ remains a distribution over panels.

Explicitly, we then have

$$\|\pi - \bar{\pi}\|_\infty = \|Mp - M\bar{p}\|_\infty \quad (\text{C.1})$$

$$= \|M(x - \bar{x})\|_\infty \quad (\text{C.2})$$

$$= \frac{1}{2m} \left\| \begin{bmatrix} M^* \\ \bar{M}^* \end{bmatrix} y \right\|_\infty, \quad (\text{C.3})$$

where $y := 2m(\bar{x} - x)$. Because ℓ_∞ is a maximum, this is

$$\geq \frac{1}{2m} \|M^* y\|_\infty \quad (\text{C.4})$$

$$= \frac{1}{2m} \left\| \frac{1}{2} H_t^* y + \frac{1}{2} \mathbf{1}^{(N-1) \times N} y \right\|_\infty \quad (\text{C.5})$$

$$= \frac{1}{4m} \|H_t^* y\|_\infty, \quad (\text{C.6})$$

where $\sum_i y_i = 0$ because we require that \bar{p} remain a distribution. Then since $y \in (2\mathbb{Z} + 1)^{N-2}$, by Lemma B.3 we have

$$\geq \frac{\sqrt{N}}{4m} \quad (\text{C.7})$$

$$= \Omega\left(\frac{\sqrt{k}}{m}\right), \quad (\text{C.8})$$

since $k = N - 1$. Again since $\bar{\mathcal{D}} \subseteq \{p + \bar{\Delta}/2m\}$, we then have

$$\min_{\bar{p} \in \bar{\mathcal{D}}} \|\pi - \bar{\pi}\|_\infty = \Omega\left(\frac{\sqrt{k}}{m}\right).$$

Since π is uniform by construction (and so these p and π are optimal with respect of Maximin), this is a lower bound on the discrepancy of each marginal which was minimal *before deviation*. It finally remains to show that this deviation happens in the downward direction, so that the minimum marginal decreases by at least this amount. Observe that by the construction of \bar{M}^* , for all \bar{p} we have $(M^* \bar{p})_i = -(M^* \bar{p})_i$. Therefore for any given \bar{p} , whichever coordinate i satisfies $|(\pi - \bar{\pi})_i| = \Omega(\sqrt{k}/m)$, there is a coordinate i' for which $(\pi - \bar{\pi})_{i'} = \Omega(\sqrt{k}/m)$. Therefore in this instance

$$\text{Maximin}(p^*) - \max_{\bar{p} \in \bar{\mathcal{D}}} \text{Maximin}(\bar{p}) = \Omega\left(\frac{\sqrt{k}}{m}\right),$$

as desired. \square

Lemma 4.1. *For NW-optimal p^* over a support of panels $\text{supp}(p^*)$, there exists a constant $\lambda \in \mathbb{R}^+$ such that, for all $P \in \text{supp}(p^*)$, $\sum_{i \in P} 1/\pi_i^* = \lambda$.*

Proof of Lemma 4.1. We can write the problem of finding the NW optimizing distribution over a fixed panel support $\mathcal{P} \subseteq \mathcal{K}$ as below on the left, where $NW^n(p)$ is equal to the geometric mean of π , the marginals implied by the panel distribution p (in contrast, in Section 2, we let $NW(p)$ be the n th root of the geometric mean—here we take the n -th power of NW to get the geometric mean NW^n). On the right, we’ve rewritten the program in standard form, where we set $f(p) = -NW^n(p)$, $h(p) = p_1 + p_2 + \dots + p_{|\mathcal{P}|} - 1$, and $g_j(p) = -p_j$. Observe that, $\forall j \in [|\mathcal{P}|]$, $\nabla h(p) = \mathbf{1}$ and $\nabla g_j(p) = -e_j$, where e_j is the vector of 0s with a 1 at index j .

$$\begin{array}{ll} \max_p NW^n(p) & \min_p f(p) \\ \|p\|_1 = 1 & h(p) = 0 \\ p_j \geq 0 \ \forall j \in [|\mathcal{P}|] & g_j(p) \leq 0 \ \forall j \in [|\mathcal{P}|] \end{array}$$

Now, let p^* be the optimal solution to this program, and $supp(p^*)$ be its support, i.e., the set of panels to which p^* assigns nonzero probability. Then, since the objective and constraints of the above program are continuously differentiable over their entire support (and thus at p^*), by the KKT condition Stationarity, there exist some constants λ and μ_j for all $j \in [|supp(p^*)|]$ (where $\mathbf{0}$ is the zero vector) such that

$$\nabla f(p^*) + \lambda \nabla h(p^*) + \sum_{j \in [|supp(p^*)|]} \mu_j \nabla g(p^*) = \mathbf{0} \implies (\nabla f(p^*))_j = \mu_j - \lambda$$

By dual feasibility and primal feasibility respectively, we have that $\mu_j, p_j \geq 0$ for all $j \in [|supp(p^*)|]$; by complementary slackness, we have that $\sum_{j \in [|supp(p^*)|]} \mu_j p_j = 0$. Thus, for all j , either $p_j = 0$, or $p_j > 0$ and $\mu_j = 0$. We have restricted $supp(p^*)$ to panels j in which $p_j > 0$, so we conclude that $\mu_j = 0$. It follows that

$$\frac{\partial NW^n(p^*)}{\partial p_j} = -(\nabla f(p^*))_j = -(\mu_j - \lambda) = \lambda \quad \forall j \in supp(p^*)$$

Finally, we can conclude the proof by expressing this partial derivative for fixed p_j (which as shown, has a constant value across all j) in terms of the marginals π :

$$\forall j \in supp(p^*), \quad \lambda = \frac{\partial NW^n(p^*)}{\partial p_j} = \sum_{i \in N} \frac{NW^n(p^*)}{\pi_i} \frac{\partial \pi_i}{\partial p_j} = \sum_{i \in P_j} \frac{NW^n(p^*)}{\pi_i} = NW^n(p^*) \left(\sum_{i \in P_j} \frac{1}{\pi_i} \right)$$

where P_j is the j th panel in $supp(p^*)$. The second equality is by the product rule, where each term of the resulting sum is equal to the derivative of π with respect to p_j multiplied by NW/π_i , the NW holding out the marginal of individual i . The third equality is by the fact that if $i \in P_j$, then $\partial \pi_i / \partial p_j = 1$; else, $\partial \pi_i / \partial p_j = 0$. \square

Lemma 4.2. *For Nash-Welfare optimal p^*, π^* , we have that $\pi_i^* \geq 1/n$ for all $i \in N$.*

Proof of Lemma 4.2. Let $X[P \ni i]$ be the indicator that a panel P contains individual i . Then,

$$\mathbb{E}_{P \sim p^*} \left[\sum_{i \in P} \frac{1}{\pi_i^*} \right] = \mathbb{E}_{P \sim p^*} \left[\sum_{i \in N} \frac{X[P \ni i]}{\pi_i^*} \right] = \sum_{i \in N} \frac{\mathbb{E}_{P \sim p^*}[X[P \ni i]]}{\pi_i^*} = \sum_{i \in N} \frac{\pi_i^*}{\pi_i^*} = n$$

By Lemma 4.1, we also have that $\mathbb{E}_{P \sim p^*} \left[\sum_{i \in P} \frac{1}{\pi_i^*} \right] = \lambda$, and thus the λ specified by Lemma 4.1 is equal to n . It follows that $\pi_i^* \geq 1/n \forall i \in N$; otherwise, we would have some panel P for which $\sum_{i \in P} \frac{1}{\pi_i^*} > n$, a contradiction. \square

Lemma 4.3. *For NW-optimal p^*, π^* , there exists a uniform lottery $\bar{p}, \bar{\pi}$ that satisfies $NW(p^*) - NW(\bar{p}) \leq k \|\pi^* - \bar{\pi}\|_\infty$.*

Proof of Lemma 4.3. Let π_{min}^* be the smallest marginal of any individual implied by the Nash-optimal distribution over panels p^* , i.e., $\pi_{min}^* = \min_{i \in N} \pi_i^*$. Then, to upper-bound the loss in NW, we assume an unattainable worst case that between p^*, π^* and a given uniform lottery $\bar{p}, \bar{\pi}$, all individuals probabilities suffer the largest loss of any marginal, $\|\pi^* - \bar{\pi}\|_\infty$. This first gives the multiplicative bound

$$NW(\bar{p}^*) \geq NW(p^*) \left(\frac{\pi_{min}^* - \|\pi^* - \bar{\pi}\|_\infty}{\pi_{min}^*} \right) = NW(p^*) \left(1 - \frac{\|\pi^* - \bar{\pi}\|_\infty}{\pi_{min}^*} \right).$$

Rearranging the above conclusion and then applying the facts that $NW(p^*) \leq k/n$ (trivially) and $\pi_{min}^* \geq 1/n$ (Lemma 4.2), we get the desired additive bound:

$$NW(p^*) - NW(\bar{p}) \leq NW(p^*) \cdot \frac{\|\pi^* - \bar{\pi}\|_\infty}{\pi_{min}^*} \leq \frac{k}{n} \cdot \frac{\|\pi^* - \bar{\pi}\|_\infty}{1/n} \leq k \|\pi^* - \bar{\pi}\|_\infty \quad \square$$

D Omitted Materials from Section 5

D.1 Algorithm Descriptions

Algorithms for calculating optimal panel distributions.

In this paper, we calculate optimal panel distributions across instances with respect to Maximin, NW, and Leximin objectives. To do this, we build on publicly-available code [HG20], which implements the column generation techniques from [FGG⁺21].

Rounding algorithms.

At a high level, the task solved by the PIPAGE and BECK-FIALA rounding algorithms in Section 5 can be thought of as rounding an input panel distribution p to some uniform lottery \bar{p} by rounding the STANDARD LP described in Section 3. However, neither of these rounding methods are used to directly round p ; rather, they are used to round a modified version p' , which transforms the task from rounding entries of p to multiples of $1/m$ to the task of rounding entries of p' to 0/1. The details of this transformation are described in the proof of Theorem 3.2 in Appendix B.

PIPAGE

We round p' exactly according to the Pipage Rounding algorithm specified in Gandhi *et al* [GKPS06]. We note that their algorithm is specified for the task of rounding bipartite graphs; we apply their methods by formulating our rounding problem as a star graph, where each of the $|\mathcal{K}|$ vertices surrounding the central vertex corresponds to a feasible panel P . Each edge from the central vertex i to a surrounding vertex P has a weight (which will ultimately be rounded to 0/1) equal to $x_{i,P} = p'_P$, the probability of drawing panel P from the modified version of the initial distribution p' . Gandhi *et al*'s degree preservation property guarantees the satisfaction of our adding up constraint $\|p'\| = \|\bar{p}'\|$.

BECK-FIALA

Our Beck-Fiala implementation is identical to the deterministic implementation specified in the proof of Lemma 9, Appendix B.4.1 of [FGGP20]. For details on the mapping of their setting to ours, see the proof of Theorem 3.2 in Appendix B.

Integer Programs.

IP-MAXIMIN

The below integer program computes a lottery $\bar{p} \in (\mathbb{Z}^+/m)^{|\mathcal{K}|}$, where the variables are y , the lower bound on any marginal probability; \bar{p} , the uniform lottery; and $\bar{\pi}$, the implied vector of marginals. The first constraint, along with the objective, result in the maximization of the minimum marginal. The second constraint imposes the relationship between the panel distribution \bar{p} and the marginals $\bar{\pi}$. The third constraint imposes that the resulting panel distribution \bar{p} will be a uniform lottery. The fourth and fifth constraints impose that \bar{p} is a valid distribution.

$$\begin{aligned}
 & \text{Maximize } y \\
 \text{s.t. } & \bar{\pi}_i \geq y & \forall i \in N \\
 & \sum_{\substack{P \in \mathcal{K}, \\ P \ni i}} \bar{p}_P = \bar{\pi}_i & \forall i \in N \\
 & m \bar{p}_P \in \mathbb{Z}^+ & \forall P \in \mathcal{K} \\
 & \sum_{P \in \mathcal{K}} \bar{p}_P = 1 \\
 & \bar{p}_P \geq 0 & \forall P \in \mathcal{K}
 \end{aligned}$$

IP-NW

This integer program is essentially the same as IP-MAXIMIN, except that instead of maximizing the lower bound on the marginals, it maximizes the geometric mean of the marginals by equivalently

maximizing the sum of their logarithms.

$$\begin{aligned}
& \text{Maximize} \sum_{i \in N} \log(\bar{\pi}_i) \\
& \text{s.t.} \sum_{\substack{P \in \mathcal{K}, \\ P \ni i}} \bar{p}_P = \bar{\pi}_i \quad \forall i \in N \\
& \quad m \bar{p}_P \in \mathbb{Z}^+ \quad \forall P \in \mathcal{K} \\
& \quad \sum_{P \in \mathcal{K}} \bar{p}_P = 1 \\
& \quad \bar{p}_P \geq 0 \quad \forall P \in \mathcal{K}
\end{aligned}$$

IP-MARGINALS

This IP takes as input some panel distribution p, π to be rounded, and minimizes the largest discrepancy of any resulting $\bar{\pi}_i$ from the corresponding π_i . Again, several of the constraints and variables are common with IP-MAXIMIN.

$$\begin{aligned}
& \text{Minimize} z \\
& \text{s.t.} |\pi_i - \bar{\pi}_i| \leq z \quad \forall i \in N \\
& \sum_{\substack{P \in \mathcal{K}, \\ P \ni i}} \bar{p}_P = \bar{\pi}_i \quad \forall i \in N \\
& \quad m \bar{p}_P \in \mathbb{Z}^+ \quad \forall P \in \mathcal{K} \\
& \quad \sum_{P \in \mathcal{K}} \bar{p}_P = 1 \\
& \quad \bar{p}_P \geq 0 \quad \forall P \in \mathcal{K}
\end{aligned}$$

D.2 Implementation Notes and Algorithm Runtimes

Our experiments were implemented in Python and run on a 13-inch MacBook Air (2018) with a 1.6 GHz Intel Core i5 processor.

Runtimes of PIPAGE, BECK-FIALA, and IP-NW on rounding an unconstrained distribution are given in the table below.⁹ We optimized IP-NW with Gurobi using its built-in piecewise linear approximation of logarithms (given that IP-NW is nonlinear) with the parameter controlling the error in the piecewise approximation set to `FuncPieceError=0.0001`. This worked quite well in most instances, getting within $1/m$ of optimal fairness on 10 out of 11 instances.

IP-MAXIMIN and IP-MARGINALS were run in Gurobi and struggled to converge completely (even after many hours), but showed good performance after a short time. The results in the paper show their solutions after 30 minutes of run-time.

* indicates capped at 7200s (2 hours). Time is measured in seconds. All times given (except those that timed out) represent the average over 3 runs.

D.3 Analysis of Nash Welfare Fairness Preservation (Figure corresponding to Figure 1)

Here we give the corresponding analysis from Figure 1 for NW. We see, first that there is some algorithm in every instance that achieves within $0.1/m$ of $NW(p^*)$, where p^* is the NW optimizing unconstrained distribution. This indicates that the cost of transparency to NW in practice is essentially 0. We note that in a few instances, IP-NW, which should theoretically dominate all other algorithms, is outperformed by either PIPAGE or BECK-FIALA. As we discuss in Appendix D.2, this is due to small errors in the integer optimization errors.

We find that our theoretical upper bounds on NW loss are less useful than those on the Maximin loss, because they are multiplied by an additional factor of k , while the value of the NW objective

⁹Clarification: Although we ran 1000 replicates of PIPAGE for our experiments, times below reflect one single run of the algorithm.

Table 2: Run-times for PIPAGE, BECK-FIALA, and IP-NW

Instance	PIPAGE	BECK-FIALA	IP-NW
sf(a)	1.5	1.6	17.1
sf(b)	1.3	1.3	27.8
sf(c)	1.0	1.1	33.1
sf(d)	2.1	2.3	40.6
sf(e)	17.0	28.3	7245*
cca	4.4	6.4	7207*
hd	1.5	1.7	120.1
mass	0.4	0.4	3.4
nexus	2.8	3.2	21.1
obf	2.3	2.4	22.3
ndem	2.2	2.6	34.8

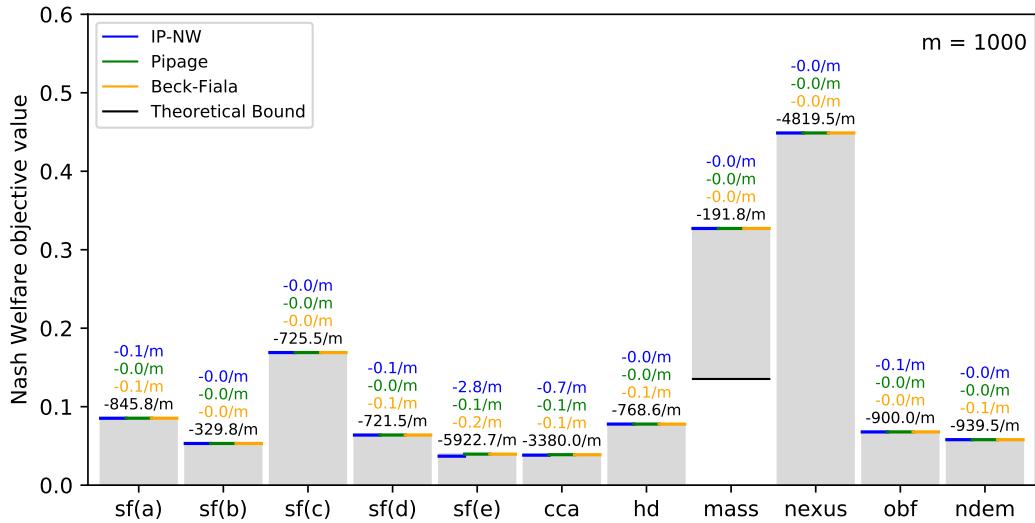


Figure 3: $m = 1000$. Shaded regions extend from $NW(p^*)$, the fairness of the optimal unconstrained distribution, down to the minimum fairness implied by the tightest theoretical upper bound in that instance (in all instances but “obf” Theorem 3.3 is tightest). Each algorithm or bound’s loss relative to $NW(p^*)$ is written above in the corresponding color. PIPAGE NW values are the average of 1000 runs (standard deviation $< 10e-17$ in all instances).

falls within a similar range to the Maximin objective. We note, however, that these bounds would be useful for larger m : currently, the maximum possible losses implied by the bounds fall between $191/m = 0.191$ and $5922/m = 5.922$. If we increased m by a factor of 100 to $m = 100,000$ (this would mean drawing 5 lottery balls instead of 3), then our bounds would be nearly tight to optimal in multiple instances (e.g., in “sf(a)”, this would yield a loss of 0.008), and would be meaningful in all instances.

D.4 Analysis of Leximin Preservation (Figures corresponding to Figure 2)

Here we give the corresponding analysis from Figure 2 for all other instances. In all instances, the conclusions we draw are essentially the same as those drawn from Figure 2: in all instances, all algorithms almost exactly preserve the Leximin-optimal marginals. Our theoretical bounds are meaningful, but we consistently outperform them in practice. Note that in all figures, PIPAGE marginals are the average of 1000 runs (maximum standard deviation of any marginal $< 10e-14$ in all instances).

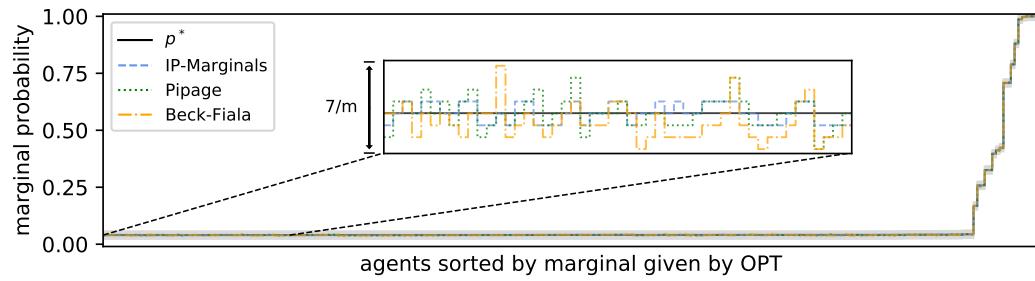


Figure 4: sf(b)

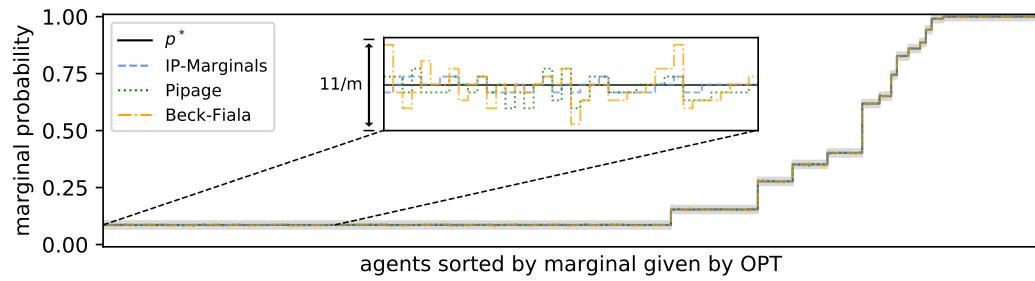


Figure 5: sf(c)

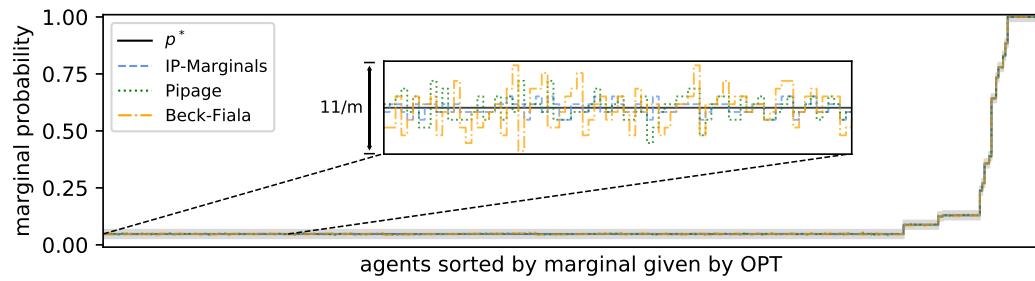


Figure 6: sf(d)

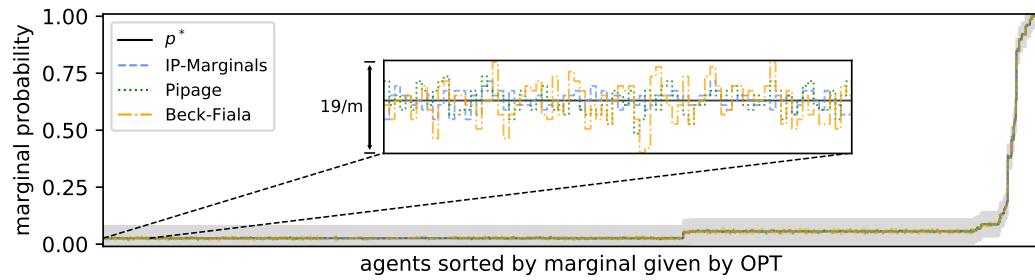


Figure 7: sf(e)

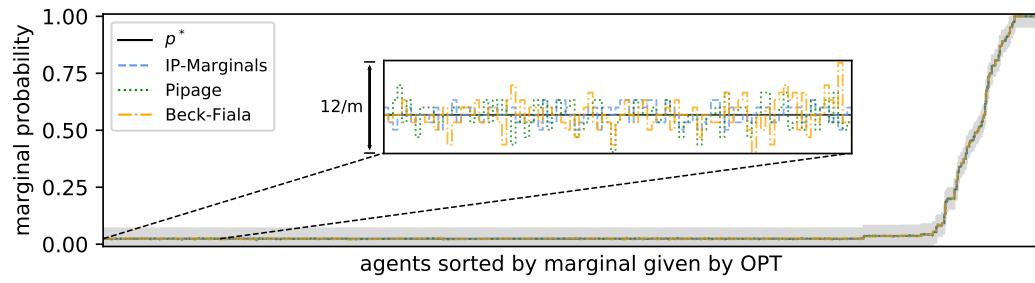


Figure 8: cca

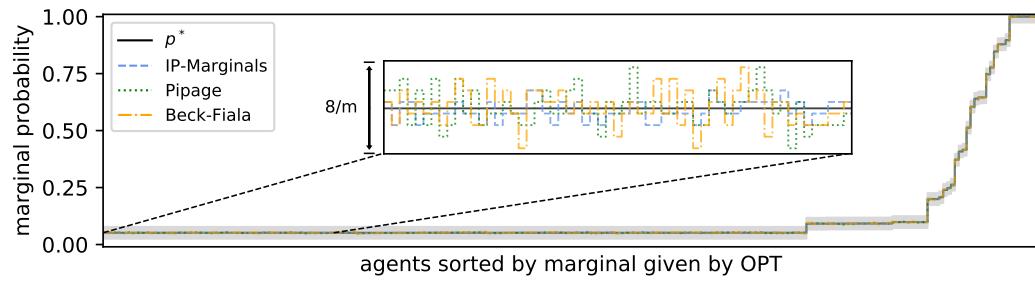


Figure 9: hd

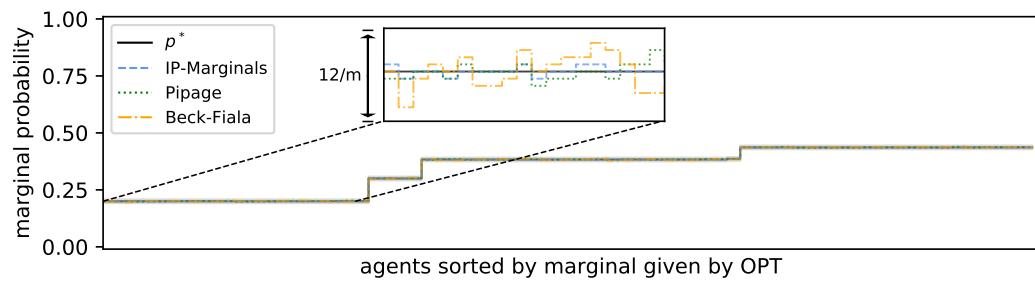


Figure 10: mass

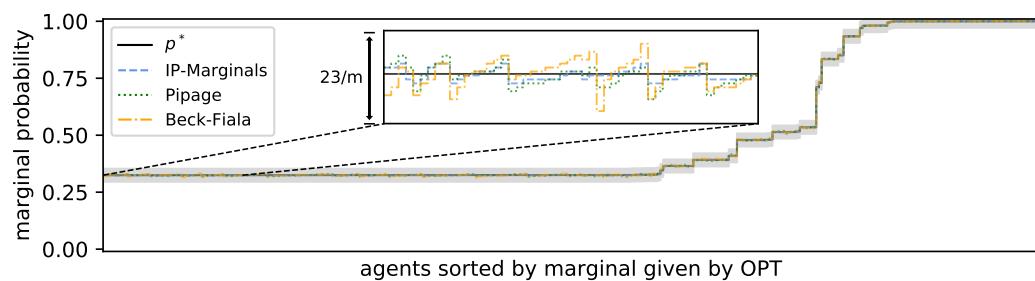


Figure 11: nexus

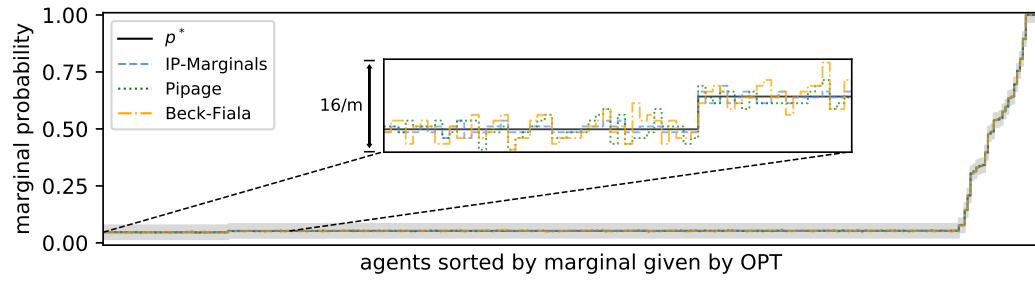


Figure 12: obf

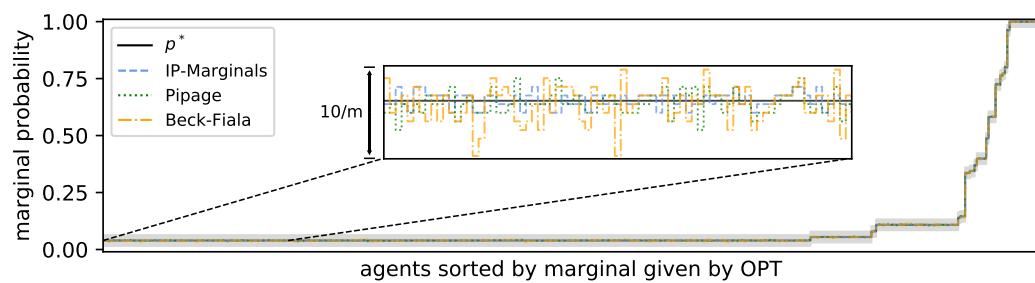


Figure 13: ndem