Price of Anarchy of Algorithmic Monoculture

Robert Kleinberg, Erald Sinanaj, and Éva Tardos

Cornell University, Ithaca NY, USA
rdk@cs.cornell.edu, erald@cs.cornell.edu, eva.tardos@cornell.edu

Abstract. Several recent works investigate the effects of *monoculture*, the ever increasing phenomenon of (possibly) self-interested actors in a society relying on one common source of advice for decision making, with an archetypal driving example being the growing adoption and predictive power of machine learning models in matching markets, e.g. in hiring. Kleinberg and Raghavan (PNAS, 2021) introduced a model that captures the effects of monoculture in a one-sided matching market with advice, demonstrating that a higher accuracy common signal (such as an algorithmic vendor) might incentivize society as a whole to rationally adopt it, but as a collective it would be better off if each instead adopted less accurate, but private advice. We generalize their model and address the open question of their work in quantifying the *social welfare loss*. We find that monoculture and more generally decentralized optimization is close to optimal: we show a tight constant bound of 2 on the price of anarchy (and more general notions) for the induced game.

1 Introduction

Algorithmic monoculture describes the now prevalent pattern of collective reliance on algorithmically curated advice. While the predictive power of algorithms and machine learning models usually outclasses traditional means, convergence and over-dependence can lead to detrimental effects in social welfare [12,15,9] and fairness concerns over outcome homogenization [17,5,2].

Our work considers a one-sided matching market under limited advice and specifically a model that is directly influenced by the one introduced by Kleinberg and Raghavan in [12] that kick-started the study of the effects of algorithmic monoculture. The motivation stemmed from the setting of *algorithmic hiring*: in their model a number of firms each seek to hire one candidate out of a pool of many, sharing a common ground truth valuation. However, they do not know *a priori* this valuation (or future performance). To the aid of their decisions, they have access to a public algorithmic vendor that relays a single shared noisy ranking to any who choose to view it and a private human advisor (employee) that also produces a ranking for their corresponding firm, independently of any other. They model both of them as ranking distributions of the same family, parameterized by an accuracy term; the algorithm is more accurate than a human advisor. Each firm chooses either the algorithm or the human for advice and they proceed in random order to hire the top remaining candidate in the ranking they received.

The theme of [12] is, as they state, a counterintuitive fact: the choice of the algorithm might be a dominant strategy for each firm, but welfare and thus average utility might be better if every firm instead used the less accurate human evaluator. They show this to be the case for two firms, under some natural assumptions, and experimentally verify that this phenomenon can arise for multiple firms as well. However, this surprising qualitative conclusion does not shed light to the degree of the negative effect that monoculture impacts. Our work is devoted to providing a measure to this welfare loss, answering their open question with a rather reassuring answer: monoculture (and more generally selfish behavior) is not that bad. We show that the *price of anarchy* for this game is 2 (which is asymptotically tight), for *any* number of agents and *any* number of "advice sources", thus bounding the welfare loss even when there is no dominant strategy (and we also extend this for general notions of equilibrium).

We show this to be the case for a generalization of their model, substantially relaxing the assumptions about the structure of the advice acquisition and the relevant ranking distributions, which allows us to apply our findings to a much broader landscape, capturing a multitude of realistic scenarios. The properties we seek for the advice sources are attained by a very rich class of distributions, which we term *Schur-concave* rankings; as an example, any ranking produced by additive i.i.d. noise falls under this class, whenever the noise distribution is *log-concave* (with two examples being the Gaussian and Laplacian, considered in [12]). We also show that the quite well known Mallow's Model rankings enjoy the same properties that lead to our bounds. Further, as mentioned, our results do not assume the existence of a dominant strategy and allow for arbitrary assymetry between the firms, specifically any number and combination of unique idiosyncratic and 'common' signals. Additionally, we extend our results to a more realistic mechanism, that allows for more expressive strategic behavior and we show that under minimal assumptions a firm does not have to *strategize* any differently from the previous model.

The paper is structured as follows. In the immediately succeeding section we discuss related work. In Section 3 we formalize specifics of our problem. In Section 4 we show our main result; a constant price of anarchy. In Section 5 we discuss the generality of our assumptions and the strategic extension. In Section 6 we establish that the our price of anarchy bound is tight and show that if the natural assumption we impose on the rankings is violated we have catastrophic results, as the price of anarchy becomes unbounded.

2 Related Work

A number of recent papers explore the theme of *algorithmic monoculture*, addressing various questions about the impacts of substituting algorithmically-derived rankings and recommendations for human preferences in strategic settings. These works typically model the algorithmically-derived preferences as more accurate (which tends to improve social welfare) but less varied (which tends to reduce social welfare), leading to subtle questions about which of these effects predominates.

This phenomenon was explored in the setting of one-sided matching by Kleinberg and Raghavan [12], in two-sided matching markets by Peng and Garg [15], and in the setting of competing classifiers by Jagadeesan, Jordan, Steinhardt, and Haghtalab [9]. While each of the aforementioned papers introduces examples of settings in which algorithmic monoculture leads to reduced social welfare, to the best of our knowledge the present work is the first to bound the worst-case social welfare loss due to algorithmic monoculture. We establish such bounds in a one-sided matching model directly inspired by [12].

The model of one-sided matching we investigate assumes that the mechanism used is Random Serial Dictatorship (also known as Random Priority). While the price of anarchy of one-sided matching mechanisms, including Random Serial Dictatorship, has been investigated in several prior works, the model we study differs in two key respects. First, it is a common value environment: any candidate has the same value to all agents. Consequently, the social welfare of a matching depends only on the set of candidates matched, rendering price-of-anarchy questions vacuous in the special case when candidates and agents are equinumerous. Second, most importantly, the agents in our model are uncertain of the candidates' values. In fact, the main strategic decision they face is not which candidate to select, but which *signal* to obtain in order to advise their ranking of the candidates.

In the classical model of matching when firms have private values for the candidates, for markets with an equal number of agents and candidates, under a standard unit-sum normalization assumption on agents' values, Christodoulou et al. [6] obtained a tight bound of $\Theta(\sqrt{n})$ on the price of anarchy of both Random Serial Dictatorship and another widely-studied mechanism, the Probabilistic Serial mechanism of Bogomolnaia and Moulin [4]. Other earlier works had focused on welfare guarantees for truthful mechanisms [1,8] and some investigated novel measures of welfare [3]. Jiang, Ndaiye, Vetta, and Wu [10] analyzed the price of anarchy of one-sided matching mechanisms in asymmetric markets where candidates may outnumber agents. In particular, they showed that the price of anarchy in markets with n agents can be asymptotically greater than $\sqrt{n}$ when the number of candidates is sufficiently large compared to n.

3 Setup

Game specification There are n firms (agents) with a shared but unknown cardinal preference over d candidates (items) and a unit demand over them. Any candidate awards a firm some non-negative utility or value. Firms have access to *ranking technologies* (advice sources). Each firm chooses one such technology, from those available to her and receives a ranking sample, not necessarily independent of one another firm receives (below we discuss this in further length). *Random Serial Dictatorship* dictates the order in which each firm is called: nature draws a uniformly random ordering and each firm is called in that order. When a firm is called, she receives the top remaining candidate from her chosen ranking and this is public knowledge. This is the setting of [12], which we will mostly focus on; we

will call this mechanism *Obedience-constrained* Random Serial Dictatorship. We will also investigate a less strict mechanism, which is identical to the previous but different in one important aspect: a firm can decide freely about her choice of a candidate when it's her turn. We will call this mechanism *Unconstrained* Random Serial Dictatorship. More discussion about the unconstrained model will follow in Section 5; unless explicitly stated, all definitions and claims will be for the former model.

Ranking technologies A ranking technology is a distribution on rankings of candidates (permutations of the set $[d] := \{1, \ldots, d\}$). There exists a *common signal space* $\mathcal{A}$ of independent ranking technologies, such that if any two firms choose the same strategy from $\mathcal{A}$ they receive the same sample. Each firm has access to some arbitrary subset $\mathcal{A}_i \subset \mathcal{A}$ and an *idiosyncratic* space of independent ranking technologies $\mathcal{H}_i$, which is mutually disjoint between firms, (that is $\mathcal{H}_i \cap \mathcal{H}_j = \emptyset$ for $i \neq j$). We note that ranking technologies are mutually independent (for any realization of the candidate values). The action space of the firms in line with their model corresponds to $\mathcal{A}$ and $\mathcal{H}_i$ being singleton sets, with $\mathcal{A}_i = \mathcal{A}$ and each idiosyncratic technology being identically (but independently) distributed. The pure strategy space is the set $\mathcal{S} := \times_{i\in[n]}\mathcal{R}_i$ (with $\mathcal{R}_i := \mathcal{A}_i \cup \mathcal{H}_i$). Further, we will refer to $\mathcal{R} := \bigcup_{i\in[n]} \mathcal{R}_i$ as the *advice space* of the game.

Valuations, welfare and notation Candidates are identified with indices $i \in [d]$ and their value $x(i)$. The joint value vector $\mathbf{x}$ is drawn from some distribution $\mathcal{X}$ supported on a subset of $[0, \infty)^d$. Whenever we have a ranking technology R we will usually denote with $r = (r_1, \ldots, r_d)$ the ranking (i.e., permutation of $[d]$) sampled from the corresponding distribution $\mathcal{F}_R(\mathbf{x})$. If C is some subset of $[d]$, r^{-C} will denote the partial ranking induced by removing the candidates in C from consideration. Thus, in line with the previously defined notation, r_i^{-C} denotes the i-th ranked candidate in the corresponding ranking (whenever it's well defined). We will usually denote a pure profile by the letter $\mathbf{s} \in \mathcal{S}$. The utility $u_i(\mathbf{s})$ of some firm in profile $\mathbf{s}$ is the expected value of the candidate she hires and the social welfare $SW(\mathbf{s})$ is the usual *utilitarian* one (the utility sum). A socially optimal profile is one that maximizes the social welfare. The Price of Anarchy [13] is the ratio of the optimal welfare to the worst-case Nash equilibrium welfare.

4 The price of anarchy

Before we state our main result we define a very natural property that we would like a ranking technology to have.

Definition 1 (Stochastically Consistent). *A ranking technology R is* stochastically consistent *iff for any two subsets M and C of $[d]$ with cardinality q and any enumeration $\{m_i\}_{i\in[q]}$, $\{c_i\}_{i\in[q]}$ of them, for any realization $\mathbf{x}$ of the joint value*

vector, the following holds for all $i, j \notin M$, $k, \ell \notin C$ *with* $(i-j)(x(k)-x(\ell)) \leq 0$*:*

$$\mathbb{P}_{r\sim\mathcal{F}_R(x)}\big((r_i, r_j) = (k, \ell) \mid (r_{m_1}, \ldots, r_{m_q}) = (c_1, \ldots, c_q)\big) \geq \mathbb{P}_{r\sim\mathcal{F}_R(x)}\big((r_i, r_j) = (\ell, k) \mid (r_{m_1}, \ldots, r_{m_q}) = (c_1, \ldots, c_q)\big).$$

That is, the ranking obeys a notion of *independence of irrelevant alternatives*: conditioned on two candidates k and ℓ being ranked at two specific spots (i and j) in any order and conditioned on the set of candidates ranked in some other arbitrary positions, it is more probable that the ranking correctly represents the relative order of the aforementioned pair, rather than not. In Section 5 we show how deceivingly broad it is; demonstrating that it captures almost any natural ranking distribution one can think of.

We extend this definition, declaring an *advice space* stochastically consistent when all its ranking technologies are stochastically consistent. Our main result is the following favorable conclusion about the social welfare loss of monoculture and more generally of any Nash equilibrium.

Theorem 1. *The Price of Anarchy of the Obedience-Constrained Random Serial Dictatorship mechanism, in which firms obtain their rankings from a stochastically consistent advice space, is bounded above by 2.*

Before presenting the full proof of the theorem, we sketch the intuition behind the result. An initial observation is that, since agents have shared (cardinal) preferences in our model, the total utility of all agents is equal to the total value of the candidates hired. The main idea of the proof is to consider two cases, depending on the total value of the candidates that each firm gets in the socially optimal profile $\mathbf{s}^*$, but would have been "taken" from them in the equilibrium profile $\mathbf{s}$ by the firms that went before her. The value of this set of candidates will be denoted in the proof by $W^{(\beta)}(\mathbf{s} \cap \mathbf{s}^*)$, where β is the order of the firms. If its expectation is at least half that of the optimum social welfare, then clearly the price of anarchy is bounded above by two. Otherwise, at least half of the expected total value of the candidates hired in the socially optimal profile must be attributable to candidates that would not be "taken" by those before the respective firm in the equilibrium profile.

Thus, in the second case, if a firm at random in the equilibrium had unilaterally deviated to her strategy in the optimal profile, she would be expected to do reasonably well by at least gaining a fraction of this "unshared" welfare, assuming that rankings are "well behaved" enough (which we discuss immediately below). The argument concludes by recalling the definition of the equilibrium; the social welfare of the equilibrium should upper bound the total expected utility of every such firm unilaterally deviating to this strategy.

The next lemma establishes the "well behaved" property of the rankings that we need for our argument to hold and shows that it is the case for our modeling assumptions. In short, we claim that *stochastic consistency* of the ranking strategies implies the following fact. For any two pure strategy profiles $\mathbf{s}$,

$\mathbf{s}^*$, conditioned on the order of firms being β, whenever the random candidate that firm i hires in profile $\mathbf{s}^*$ is not selected by the firms before her in $\mathbf{s}$, she would have done just as well or better, if she was using s_i^* instead in profile $\mathbf{s}$. While this idea is intuitive, it turns out to not be the case in general: in some cases a higher ranked candidate being available may correlate with the value of the candidate being poor in expectation. We show that our intuition is true however, when the advice space is stochastically consistent. The proof of the statement below can be found in the Appendix.

Lemma 1. *Consider an Obedience-Constrained Random Serial Dictatorship mechanism with a stochastically consistent advice space. Let $c_i^{(\beta)}(\mathbf{s})$ denote the candidate hired by firm i in profile $\mathbf{s}$ when the order of firms is β. Let also $x_i^{(\beta)}(\mathbf{s}) := x(c_i^{(\beta)}(\mathbf{s}))$ be the value of this candidate. Further, define $C_i^{(\beta)}(\mathbf{s})$ to be the set of candidates hired by the firms going before i in profile s.*

For any two pure strategy profiles s, s^* the below holds:

$$\mathbb{E}\left[x_i^{(\beta)}(s_i^*, \mathbf{s}_{-i}) - x_i^{(\beta)}(\mathbf{s}^*) \;\middle|\; \mathbf{x}, \beta, c_i^{(\beta)}(\mathbf{s}^*) \notin C_i^{(\beta)}(\mathbf{s})\right] \geq 0$$

where the expectation is over the randomness of the samples drawn from the ranking technologies and $\mathbf{x}$ represents the joint value vector of the candidates.

Now we are ready to prove our main result.

Proof (Theorem 1). Consider a pure strategy profile $s = (s_1, \ldots, s_n)$. The outcome of the selected candidates depends on the rankings sampled from the technologies adopted by the agents, but also on the order β that Random Serial Dictatorship draws (and the random candidate values $\mathbf{x}$). Let the socially optimal profile be s^* and define the following quantity informally discussed above:

$$W^{(\beta)}(\mathbf{s} \cap \mathbf{s}^*) := \mathbb{E}\left[\sum_{i\in[n]} x_i^{(\beta)}(\mathbf{s}^*) \cdot \mathbb{1}\left(c_i^{(\beta)}(\mathbf{s}^*) \in C_i^{(\beta)}(\mathbf{s})\right) \;\middle|\; \text{RSD order } = \beta\right]$$

(Note that when $\beta(i) = 1$ we have $C_i^{(\beta)}(\mathbf{s}) = \emptyset$ and thus these terms are zero.) Similarly, define:

$$W^{(\beta)}(\mathbf{s}^* \setminus \mathbf{s}) := \mathbb{E}\left[\sum_{i\in[n]} x_i^{(\beta)}(\mathbf{s}^*) \cdot \mathbb{1}\left(c_i^{(\beta)}(\mathbf{s}^*) \notin C_i^{(\beta)}(\mathbf{s})\right) \;\middle|\; \text{RSD order } = \beta\right]$$

Now, we distinguish two cases.

Case 1: When $\underline{\mathbb{E}_\beta\left[W^{(\beta)}(\mathbf{s} \cap \mathbf{s}^*)\right] \geq 1/2 \cdot SW(\mathbf{s}^*)}$:

In this case, we have $SW(\mathbf{s}) \geq \mathbb{E}_\beta\left[W^{(\beta)}(\mathbf{s} \cap \mathbf{s}^*)\right] \geq 1/2 \cdot SW(\mathbf{s}^*)$, as claimed. The first inequality holds because $SW(\mathbf{s})$ represents the expected combined value of *all* candidates hired in profile $\mathbf{s}$, whereas $\mathbb{E}_\beta\left[W^{(\beta)}(\mathbf{s} \cap \mathbf{s}^*)\right]$ represents the expected combined value of a subset of the candidates hired in profile $\mathbf{s}$.

Case 2: When $\underline{\mathbb{E}_\beta\left[W^{(\beta)}(\mathbf{s}\cap\mathbf{s}^*)\right] < 1/2\cdot SW(\mathbf{s}^*)}$:

In this case, we have that equivalently: $\mathbb{E}_\beta\left[W^{(\beta)}(\mathbf{s}^*\setminus\mathbf{s})\right] > 1/2\cdot SW(\mathbf{s}^*)$. Now, we use a proof outline reminiscent of the classical smoothness [16] framework for price of anarchy. Suppose a firm in the equilibrium unilaterally deviates to the strategy she employs in the optimal profile $\mathbf{s}^*$. The idea is to observe that if her choice in the optimal profile would not be taken by those before her in the equilibrium, she would hire either the candidate she would get in the optimal profile, or someone higher up in her list. In the latter case, we claim that stochastic consistency (of the advice space) guarantees that her expected utility can only increase that way.

More formally, we have that:

$$\begin{aligned}
\sum_{i\in[n]} u_i(s_i^*,\mathbf{s}_{-i}) &= \mathbb{E}\left[\sum_{i\in[n]} x_i^{(\beta)}(s_i^*,\mathbf{s}_{-i})\right]\\
&\geq \mathbb{E}\left[\sum_{i\in[n]} x_i^{(\beta)}(s_i^*,\mathbf{s}_{-i})\cdot\mathbb{1}\left(\mathcal{E}_i\right)\right], \text{ def. } \mathcal{E}_i := \left\{c_i^{(\beta)}(\mathbf{s}^*)\notin C_i^{(\beta)}(\mathbf{s})\right\}\\
&= \sum_{i\in[n]} \mathbb{E}_{\mathbf{x},\beta}\mathbb{E}\left[x_i^{(\beta)}(s_i^*,\mathbf{s}_{-i})\cdot\mathbb{1}\left(\mathcal{E}_i\right)\,\middle|\,\mathbf{x},\beta\right]\\
&= \sum_{i\in[n]} \mathbb{E}_{\mathbf{x},\beta}\Big[\mathbb{P}\left(\mathcal{E}_i\mid\mathbf{x},\beta\right)\mathbb{E}\left[x_i^{(\beta)}(s_i^*,\mathbf{s}_{-i})\mid\mathbf{x},\beta,\mathcal{E}_i\right]\Big]
\end{aligned}$$

And now notice that by Lemma 1 this is:

$$\begin{aligned}
&\geq \sum_{i\in[n]} \mathbb{E}_{\mathbf{x},\beta}\Big[\mathbb{P}\left(\mathcal{E}_i\mid\mathbf{x},\beta\right)\mathbb{E}\left[x_i^{(\beta)}(\mathbf{s}^*)\mid\mathbf{x},\beta,\mathcal{E}_i\right]\Big]\\
&= \mathbb{E}_\beta\left[W^{(\beta)}(\mathbf{s}^*\setminus\mathbf{s})\right]\\
&> 1/2\cdot SW(\mathbf{s}^*)
\end{aligned}$$

Now assume $\mathbf{s}$ is a pure equilibrium, then either case 1 holds and so we immediately conclude that $SW(\mathbf{s})\geq 1/2\cdot SW(\mathbf{s}^*)$ or not, in which case observe that by the equilibrium condition:

$$SW(s)\geq\sum_{i\in[n]} u_i(s_i^*,\mathbf{s}_{-i}) \underset{\text{Case 2}}{\geq} 1/2\cdot SW(\mathbf{s}^*)$$

Finally, notice that nothing forbids us to extend this to mixed equilibria. We just need to consider each case in expectation over the randomness of the mixed strategy equilibrium, meaning, now we define $W^{(\beta)}(\sigma\cap\mathbf{s}^*) := \mathbb{E}_{\mathbf{s}\sim\sigma}\left[W^{(\beta)}(\mathbf{s}\cap\mathbf{s}^*)\right]$ for a mixed strategy profile σ. The proof then follows the same recipe, distinguishing cases according to what proportion of the optimal welfare this quantity represents. In fact, we can extend this in the same fashion to other general

equilibrium concepts (e.g. correlated equilibria) or no regret sequences. This is because at the core of our proof (i.e., Case 2) lies a smoothness-type argument which allows us to express this stronger guarantee of "Robust Price of Anarchy" [16]. □

5 Modeling assumptions and extension

In this section we discuss and establish the generality of our assumptions.

First of all, we demonstrate that the most natural and widely considered ranking distributions do indeed satisfy stochastic consistency. We show that it holds for the well-known *Mallow's Model* distribution and for a very rich class of ranking technologies that we term *Schur-concave rankings*.

Mallow's model [14] has a long history in the literature of decision theory. It dictates that the probability of observing a certain ranking π is proportional to $\phi^{d(\pi)}$, where d is a distance from the true ranking and $\phi \in (0,1)$ is a *dispersion* parameter. One of the most frequently used distance notions is the Kendall-Tau distance [11], which measures the number of inversions (incorrectly ranked pairs) of π from the true ranking.

Theorem 2. *The Mallow's Model of rankings with the Kendall-Tau distance is* Stochastically Consistent.

Proof. That Mallow's Model is *stochastically consistent* is fairly immediate: any realization where some pair is ranked correctly can be mapped to one where they switch positions (and obviously vice-versa). This introduces at *least* one inversion, the one between the pair. Therefore, witnessing this realization instead is less than $\phi < 1$ times as likely, than the one where the relative ordering of the pair was correct. A standard argument using the law of total probability then implies the first claim. □

Before we define Schur-concave rankings we need the notion of an *Additive noise ranking*, which ranks candidates by noisy score estimation. In such a ranking distribution, first, a joint noise vector $\boldsymbol{\varepsilon}$ is sampled and then, candidates are ranked in decreasing order of their scores: $x_i + \varepsilon_i$. The below definitions serve for defining *Schur-concave rankings*:

Definition 2 (Majorization). *A vector $\mathbf{x} \in \mathbb{R}^n$ majorizes another vector $\mathbf{y} \in \mathbb{R}^n$ (and we symbolize $\mathbf{x} \succ \mathbf{y}$) whenever we have that:*

$$\sum_{i=1}^{k} x_i^{\downarrow} \geq \sum_{i=1}^{k} y_i^{\downarrow} \text{ for all } k \in [n]$$

and further: $\sum_{i=1}^{n} x_i = \sum_{i=1}^{n} y_i$, where above $x_i^{\downarrow}$ is the i-th largest entry of $\mathbf{x}$.

Definition 3 (Schur-Concavity). *A function $f : \mathbb{R}^n \to \mathbb{R}$ is said to be* Schur-Concave *whenever for all pairs $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ such that $\mathbf{x} \succ \mathbf{y}$:*

$$f(\mathbf{x}) \leq f(\mathbf{y})$$

A vector that is majorized by another is more "spread out" than the latter and so a first intuition for Schur-concave distributions is that they assign less probability mass to more 'extreme' points. With this information we are now ready to state our definition.

Definition 4 (Schur-Concave Ranking). *We call a ranking technology* Schur-Concave *if the distribution of the ranking can be modeled by an additive noise ranking, such that the joint density of the noise vector is Schur-Concave.*

The below claim is a quite well-known fact in majorization theory (see e.g. [7] and notice that this is a corollary of *Fact* 13.24 and *Fact* 13.29).

Proposition 1. *Additive noise rankings with identically and independently distributed log-concave noise (in each candidate) form a* Schur-concave *ranking.*

Some examples of log-concave distributions are possibly the most well known: the Gaussian distribution, the Laplacian, the uniform distribution and many more. The below theorem establishes the breadth of the class of distributions that our notion of "well-behavedness" (stochastic consistency) encompasses. We state its proof in the Appendix.

Theorem 3. Schur-concave rankings *are* Stochastically Consistent.

In the rest of this section we investigate the constraint that the mechanism imposes on the firms (and its relaxation). Specifically, it is assumed that when it's her turn, each firm follows her ranking sample, in choosing the top ranked available candidate, no matter what she knows about the state of the world. However, this is not generally true for rational actors, if we (naturally) assume that they have the choice of selecting from the remaining candidates however they wish. This persists even if they don't have a Bayesian belief about the rest of the firms (for example they might have known that all firms have access only to well-behaved algorithmic vendors), depending on the candidate distribution and thus one might question the assumption that a firm will trust her ranking "blindly".

This behavior is also a problem in that it complicates a firm's decision making: it would be simpler if one was reassured that the best she could do in expectation, at any point, is to follow the ranking she sampled. The statement below shows that a very simple assumption about the world makes this desideratum (which we will call *incentive compatibility*) true. Under this assumption, even if a firm knew exactly what technology the ones before her used, not just what candidate they received, she is better off choosing her top ranked choice that is available in the sample she received. Below we define the assumption in question.

Definition 5 (Permutation Invariant Candidate Distribution). *A candidate distribution $\mathcal{X}$ is* permutation invariant *iff for any permutation π of the set $[d]$ and any candidate value vector $\mathbf{x} \in [0, \infty)^d$ the following holds for the probability density function concerning the joint values:*

$$f_{\mathcal{X}}(\mathbf{x}) = f_{\mathcal{X}}(\pi \circ \mathbf{x})$$

where $\pi \circ \mathbf{x}$ denotes the permutation of the components of $\mathbf{x}$ according to π.

One can readily verify that this is satisfied by any product distribution with identical univariates (case of i.i.d. values). More generally, the condition does not require mutual independence of the candidate values, merely a notion of unidentifiability or symmetry in the *likelihood* of the realizations of the world: ie, a name, tag, or other irrelevant attribute should not predict one's potential.

We claim the following, which asserts that under this assumption, firms in the Unconstrained Random Serial Dictatorship mechanism have no incentive to behave differently than they would have in the Obedience-Constrained Random Serial Dictatorship. The proof can be found in the Appendix:

Proposition 2. *Unconstrained Random Serial Dictatorship, equipped with a* stochastically consistent *advice space is* incentive compatible *if the candidate distribution is* permutation invariant*, even if a firm knows which technologies each firm before her selected.*

The main argument is that even conditioned on some specific subset of candidates being available, the firm cannot "distinguish" any one of them as having lesser expected value from any other and thus will not gain in expectation by strategically permuting a sample of a stochastically consistent ranking (even adaptively to which candidate is recommended by her chosen ranking).

Proposition 2 implies that the same price of anarchy bound holds for the *Unconstrained* Random Serial Dictatorship, assuming that firms follow the (weakly) dominant strategy of trusting their ranking. Moreover, under some additional mild assumptions about the candidate distribution (assuming the distribution is *atomless*) and the advice space (assuming a *strict* consistency: a strict inequality for candidates with different values) trusting one's ranking is *strictly* dominant.

6 Lower bounds and violation of assumptions

In this section we show that the bound on the price of anarchy we presented under the considered model assumptions is asymptotically tight, up to $O(1/n)$ terms. We conclude the section by illustrating the almost essential necessity of considering ranking strategies that are *stochastically consistent*, as the most general but well-behaved ranking distributions. We show that when one or more ranking strategies are not "well-behaved", that is, when they do not adhere to the definition of *stochastic consistency*, there can be pathological scenarios where even a dominant strategy equilibrium can be vastly worse than the optimal social welfare, by a factor linear in the number of firms.

Theorem 4. *The Price of Anarchy bound of* 2 *is asymptotically tight.*

Proof. Consider a symmetric game with n firms and two ranking technologies for each one: common ranking technologies A and A'. Further, consider the following permutation invariant distribution on the candidate values; there are 2^n candidates, a random pair of them have values 1 and $v(n) \in (0, 1)$ (to be determined below), while all other candidates have value 0.

The ranking strategy A produces rankings that always put the candidate with value 1 on the first position; the rest of the ranking constitutes a uniform permutation on the remaining candidates. On the other hand, A' always puts the two non-zero value candidates in the first n positions, uniformly at random and the rest of the ranking also constitutes a uniformly random permutation on the remaining items.

Notice that this game instance satisfies the assumptions that lead to our bound; the ranking strategies are stochastically consistent (also the candidate distribution is permutation invariant which implies incentive compatibility for the relaxed model).

Observe that the utility of a firm choosing A, while everyone else also chooses A is:

$$u(A, A^{n-1}) = 1/n + v(n) \cdot \Theta(2^{-n})$$

Now, consider a deviation of one such firm to strategy A'. With probability $1/n$ the firm chooses first and her conditional expected utility is $1/n(1 + v(n))$. With the remaining $1 - 1/n$ probability the firm is not choosing first and cannot hire the candidate with value 1, but still hopes to get the candidate with value $v(n)$. This either happens because the candidate with value $v(n)$ ranks above all candidates with value 0 (probability $\frac{1}{n-1}$) and none of the earlier firms using strategy A picked that candidate (probability $1 - O(n \cdot 2^{-n})$) or because the candidate with value $v(n)$ is ranked below at least one candidate with value 0, but all higher-ranked candidates were chosen by earlier firms using strategy A (probability $O(n \cdot 2^{-n})$). In total, then, the expected utility of using strategy A' against profile A^{n-1} is

$$\begin{aligned} u(A', A^{n-1}) &= \tfrac{1}{n} \cdot \tfrac{1}{n} (1 + v(n)) + \left(1 - \tfrac{1}{n}\right) \left[\tfrac{1}{n-1}(1 - O(n2^{-n})) + O(n2^{-n})\right] v(n) \\ &\leq \tfrac{v(n)}{n} + \tfrac{1}{n^2} + v(n)(\tfrac{1}{n^2} + O(n2^{-n})) \end{aligned}$$

We can choose $v(n)$ so that A^n is an equilibrium, to that end we want $u(A, A^{n-1}) > u(A', A^{n-1})$ to hold and our goal is to maximize $1 + v(n)$, the asymptotic ratio of $SW(A'^n)/SW(A^n)$.

If we ignore the asymptotics, we ask to maximize $v(n)$ subject to the constraint $1/n > v(n)(1/n + 1/n^2) + 1/n^2$. Thus, for some $v(n) = 1 - 2/(n+1) - o(1/n)$ we have that A^n is indeed an equilibrium and the price of anarchy is $1 + v(n) \to 2$. □

The next example, while unrealistic, shows that in the setting of [12] even a dominant strategy equilibrium can be arbitrarily bad, when ranking strategies are not stochastically consistent.

Proposition 3. *The price of anarchy is $\Omega(n)$ (linear in the number of firms), for an arbitrary advice space, in the Obedience-Constrained Random Serial Dictatorship. Further, there exist game instances where there is a (strict) dominant strategy, the socially optimal profile has no firm using the dominant strategy and the optimal welfare is a factor of $\Theta(n)$ larger than that of the dominant strategy equilibrium.*

Proof. We consider n firms as usual and $2n$ candidates. The candidate distribution is such that uniformly at random, a set of $n-1$ of them have value 0, one of them has value n and the n rest have unique values from the increasing sequence $\{a_i\}_{i\in[n]}$ with $a_j = 1 + \frac{j(j-1)}{n-j+1} - \frac{(j-1)(j-2)}{n-j+2} - \varepsilon/n^2$ for some small $\varepsilon > 0$. Notice that this distribution is permutation invariant.

Firms have access to two ranking strategies, a common A and a private H (of the same distribution for any firm). Rankings from A always rank the candidate with value n at the first position, the $n-1$ candidates with value 0 in the next positions and then everyone else arbitrarily. On the other hand, rankings from H rank the candidates with values from the sequence in the first n positions in increasing order and everyone else arbitrarily in the next positions.

For brevity, let $A^k H^{n-k}$ denote the profile where k firms have chosen A and $n-k$ chose their private H. Clearly, $u(A, A^{n-k-1}H^k) = 1 + \frac{k}{n-k}$ by construction of A: a firm choosing A gets non-zero value (n) if it is the first to make a selection among the $n-k$ firms choosing A. Now, also observe that $u(H, A^{n-k-1}H^k) = \frac{1}{k+1}\sum_{i=1}^{k+1} a_i = 1 + \frac{k}{n-k} - \varepsilon/n^2$. Thus, A is indeed a dominant strategy.

It's clear that the socially optimal profile is H^n (changing any firm to employ A incurs a loss for non-trivial n, as the last term of the sequence is of order n^2) with a social welfare of $\frac{n(n+1)}{2} - \varepsilon/n$, while the dominant strategy equilibrium welfare is just n. □

7 Conclusion

Our work offers a tight bound on the social welfare loss in a hiring market with advice (something that can arise due to algorithmic monoculture or general selfish behavior) and identifies conditions and assumptions one should impose for such settings to be well behaved. It is an interesting avenue to investigate more complex settings, for example in regards to an even more general advice structure. The phenomenon of algorithmic monoculture in markets that can be modeled by the setting we investigate seems to not pose a grave risk to the overall welfare. Arguing about how detrimental monoculture is in regards to the price of anarchy or social welfare loss in the models of [9,15], as well as other models of algorithmic monoculture, remains an enticing question for future work.

A Missing proofs

Theorem 3. Schur-concave rankings *are* Stochastically Consistent.

Proof. Let $f : \mathbb{R}^d \to \mathbb{R}$ be the pdf of the additive noise vector. We fix arbitrary i, j, k, ℓ, M, C and their enumerations, as in the statement of Definition 1 and condition on some realization of the candidate joint value vector x.

Let r be the random ranking of the candidates induced by the additive noise scoring procedure and let $\mathcal{E}_{ij}$ denote the event $\{(r_i^{-C}, r_j^{-C}) = (k, \ell)\}$ and $\mathcal{E}$ denote the event $(r_{m_1}, \ldots, r_{m_q}) = (c_1, \ldots, c_q)$. Further, let M_{ij} be the set of joint noise

vectors $\boldsymbol{\varepsilon}$ that realize event $\mathcal{E}_{ij} \wedge \mathcal{E}$. Define $\mathcal{E}_{ji}$ and M_{ji} accordingly. Assume without loss of generality, that $i < j$ and $x(k) \geq x(\ell)$. We want to show that the following difference is non-negative:

$$\mathbb{P}\left(\mathcal{E}_{ij} \wedge \mathcal{E}\right) - \mathbb{P}\left(\mathcal{E}_{ji} \wedge \mathcal{E}\right) = \int_{M_{ij}} f(\boldsymbol{\varepsilon})\, d\boldsymbol{\varepsilon} - \int_{M_{ji}} f(\boldsymbol{\varepsilon})\, d\boldsymbol{\varepsilon}$$

Let $\delta_{k\ell} = x(k) - x(\ell) \geq 0$ and define the below bijection:

$$\xi : \mathbb{R}^d \ni \boldsymbol{y} \mapsto \boldsymbol{y} + \delta_{k\ell}(\boldsymbol{e}_k - \boldsymbol{e}_\ell) \in \mathbb{R}^d$$

Further let B be the matrix that swaps coordinates k, ℓ and define the bijection $t = \xi \circ B$. Observe that t maps elements of M_{ij} to M_{ji} and vice-versa: rankings remain the same, except items k and ℓ are swapped. To see this, let $\boldsymbol{\varepsilon} \in M_{ij}$, with $\varepsilon(m)$ the additive noise to candidate m. After the transformation the noise vector remains identical, except the k-th and ℓ-th component that would now be $\varepsilon(\ell) - \delta_{k\ell}$ and $\varepsilon(k) + \delta_{k\ell}$, resulting in scores $x(\ell) + \varepsilon(\ell)$ and $x(k) + \varepsilon(k)$ for items k and ℓ respectively (the scores of the rest of the candidates remain unchanged). Thus, the only difference in this new ranking is that k and ℓ are swapped. Now, clearly t is Lebesgue measure preserving and since it is a bijection between M_{ij} and M_{ji} it's the case that: $t^{-1}(M_{ji}) = M_{ij}$.

We have then, that:

$$\begin{aligned}\mathbb{P}\left(\mathcal{E}_{ij} \wedge \mathcal{E}\right) - \mathbb{P}\left(\mathcal{E}_{ji} \wedge \mathcal{E}\right) &= \int_{M_{ij}} f(\boldsymbol{\varepsilon})\, d\boldsymbol{\varepsilon} - \int_{t^{-1}(M_{ji})} f(t(\boldsymbol{\varepsilon}))\, d\boldsymbol{\varepsilon} \\ &= \int_{M_{ij}} \big(f(\boldsymbol{\varepsilon}) - f(t(\boldsymbol{\varepsilon}))\big)\, d\boldsymbol{\varepsilon} \qquad (1)\end{aligned}$$

Now, observe that if $\boldsymbol{z} = t(\boldsymbol{\varepsilon})$, then $z(r) = \varepsilon(r)$ for all $r \in [d] \setminus \{i, j\}$, but $z(k) = \varepsilon(\ell) - \delta_{k\ell}$ and $z(\ell) = \varepsilon(k) + \delta_{k\ell}$, as we noted before. Recall then that, since $\boldsymbol{\varepsilon} \in M_{ij}$ we must have that k is ranked above ℓ:

$$x(k) + \varepsilon(k) \geq x(\ell) + \varepsilon(\ell)$$

Which implies:

$$\varepsilon(\ell) - \delta_{k\ell} \leq \min\{\varepsilon(k), \varepsilon(\ell)\}.$$

By the same observation we also conclude that:

$$\varepsilon(k) + \delta_{k\ell} \geq \max\{\varepsilon(k), \varepsilon(\ell)\}$$

Thus, $t(\boldsymbol{\varepsilon})$ majorizes $\boldsymbol{\varepsilon}$ and by Schur-concavity of f we have that the intergrands in Equation (1) are always non-negative, leading to $\mathbb{P}\left(\mathcal{E}_{ij} \wedge \mathcal{E}\right) - \mathbb{P}\left(\mathcal{E}_{ji} \wedge \mathcal{E}\right) \geq 0$, which implies what's desired. □

Proposition 2. *Unconstrained Random Serial Dictatorship, equipped with a* stochastically consistent *advice space is* incentive compatible *if the candidate distribution is* permutation invariant*, even if a firm knows which technologies each firm before her selected.*

Proof. We will show that no firm has any incentive to strategize and ask for a candidate other than the highest ranked that is available from her sample, even if she knew not just the candidates hired before her, but also the technologies used. Suppose then she has a strategy to permute her sampled ranking, based on what she witnessed (technologies used and candidates hired) and her ranking sample.

Let $\boldsymbol{s} := (s_1, \ldots, s_k)$ for some $k \geq 0$ denote the arbitrary strategies of the firms going before her in order and let $D(\mathbf{s})$ denote the (vector of) candidates that each hired. First, consider the event where the firms hire some arbitrary candidates in that order, i.e. say $D(\mathbf{s}) = (c_1, \ldots, c_k)$. We claim that the joint distribution of the value of candidates $C' = [d] \setminus \{c_1, \ldots, c_k\}$ conditioned on $D(\mathbf{s}) = (c_1, \ldots, c_k)$ remains permutation invariant. This is because the event $D(\mathbf{s}) = (c_1, \ldots, c_k)$ is (trivially) *symmetric* with respect to these candidates and so permutation invariance of the original joint value vector, shall imply permutation invariance for the joint value vector of the candidates in C' too, conditioned on this event.

This shows that when a firm witnesses the event $D(\mathbf{s}) = (c_1, \ldots, c_k)$ she doesn't have any incentive to strategize with respect to the identities of the candidates she sees in the partial ranking r^{-D}. Thus, her strategy can as well depend only on the event $D(\mathbf{s}) = (c_1, \ldots, c_k)$ by possibly choosing some later *position* in r^{-D} than the first (and shall not depend on the sample r she sees). Say then, whenever she witnesses $D(\mathbf{s}) = (c_1, \ldots, c_k)$ she chooses the q-th ranked candidate from her list of available candidates for some $q > 1$, ie she seeks to hire r_q^{-D} instead of r_1^{-D}. The event $D(\mathbf{s}) = (c_1, \ldots, c_k)$ implies that for some $m \geq 0$ and some suitable indices ℓ_i we have that $(r_1, \ldots, r_m) = (c_{\ell_1}, \ldots, c_{\ell_m})$ ($m > 0$ implies that this ranking is sampled from an algorithmic vendor and $\mathbf{s}$ has some firms employing the same advice). But now notice the following, because of the independence of distinct technologies:

$$\begin{aligned} &\mathbb{E}\left[x(r_1^{-D}) - x(r_q^{-D}) \mid D(\mathbf{s}) = D\right] \\ &= \mathbb{E}\left[x(r_1^{-D}) - x(r_q^{-D}) \mid (r_1, \ldots, r_m) = (c_{\ell_1}, \ldots, c_{\ell_m})\right] \end{aligned}$$

And now it's not hard to see that stochastic consistency ensures that:

$$\mathbb{E}\left[x(r_1^{-D}) - x(r_q^{-D}) \mid (r_1, \ldots, r_m) = (c_{\ell_1}, \ldots, c_{\ell_m})\right] \geq 0$$

This is because:

$$\begin{aligned} &\mathbb{E}\left[x(r_1^{-D}) - x(r_q^{-D}) \mid (r_1, \ldots, r_m) = (c_{\ell_1}, \ldots, c_{\ell_m})\right] \\ &= 1/2 \cdot \mathbb{E} \sum_{\substack{k,\ell \in C' \\ k \neq \ell}} \sum_{\substack{i,j \in [d] \\ i<j}} (x(k) - x(\ell)) \cdot \Big\{ \mathbb{P}\left((r_i, r_j) = (k, \ell) \mid \mathcal{E}_D \wedge \mathcal{E}\right) \\ &\qquad - \mathbb{P}\left((r_i, r_j) = (\ell, k) \mid \mathcal{E}_D \wedge \mathcal{E}\right) \Big\} \end{aligned}$$

Where $\mathcal{E}_D$ represents the (random) rankings of the candidates in D, that are consistent with $\{r_1^{-D}, r_q^{-D}\} = \{k, \ell\} = \{r_i, r_j\}$ and $\mathcal{E}$, which is the event $(r_1, \ldots, r_m) = (c_{\ell_1}, \ldots, c_{\ell_m})$. Since the statement of stochastic consistency holds for any arbitrary conditioning on the ranking of a set of candidates (disjoint from those of interest), the above quantity must be non-negative.

This concludes our proof (as the conditioning $D(S) = D$ was arbitrary). □

Lemma 1. *Consider an Obedience-Constrained Random Serial Dictatorship mechanism with a stochastically consistent advice space. Let $c_i^{(\beta)}(\mathbf{s})$ denote the candidate hired by firm i in profile $\mathbf{s}$ when the order of firms is β. Let also $x_i^{(\beta)}(\mathbf{s}) := x(c_i^{(\beta)}(\mathbf{s}))$ be the value of this candidate. Further, define $C_i^{(\beta)}(\mathbf{s})$ to be the set of candidates hired by the firms going before i in profile s.*

For any two pure strategy profiles s, s^* the below holds:

$$\mathbb{E}\left[x_i^{(\beta)}(s_i^*, \mathbf{s}_{-i}) - x_i^{(\beta)}(\mathbf{s}^*) \;\middle|\; \mathbf{x}, \beta, c_i^{(\beta)}(\mathbf{s}^*) \notin C_i^{(\beta)}(\mathbf{s})\right] \geq 0$$

where the expectation is over the randomness of the samples drawn from the ranking technologies and $\mathbf{x}$ represents the joint value vector of the candidates.

Proof. Consider an ordering of firms β and a realization $\mathbf{x}$ of the values. We will show that this inequality is true for any choice of these. This is obviously true when $\beta(i) = 1$, i.e. when firm i chooses first. Thus we assume $\beta(i) > 1$, so then i in both profiles has some non-zero number of firms going before her.

Let's start by assuming s_i^* is a private ranking. When deviating to choosing s_i^* in profile $\mathbf{s}$ the firm will get the same candidate as in profile $\mathbf{s}^*$ or someone ranked higher (as per the conditioned event), so intuition says she should not do worse in expectation. However, this is not true in general if i is up against firms that use arbitrary (not stochastically consistent) technologies, as the event that i has a higher ranked candidate available can be non-trivially correlated with the value of that candidate, see Proposition 4 in the Appendix for an example.

Condition on the set of candidates that were hired by the firms going before i in profile $\mathbf{s}$, e.g. say $C_i^{(\beta)}(\mathbf{s}) = C$ for some arbitrary C. Consider the ranking sample that i draws using s_i^*, let it be r. When deviating to this strategy in profile $\mathbf{s}$, then, i will select some candidate r_1^{-C}. Let $\mathcal{E}_C$ denote the event that the vector of candidate values is $\mathbf{x}$, the ordering of firms is β, the set $C_i^{(\beta)}(\mathbf{s})$ is equal to C, and $c_i^{(\beta)}(\mathbf{s}^*) \notin C$. The event $\mathcal{E}_C$ can be partitioned into mutually exclusive events $\mathcal{E}_{j,C}$, where j ranges from 1 to $d - |C|$ and $\mathcal{E}_{j,C}$ represents the subset of sample points in $\mathcal{E}_C$ that satisfy $c_i^{(\beta)}(\mathbf{s}^*) = r_j^{-C}$. Conditional on event $\mathcal{E}_{1,C}$ we have $c_i^{(\beta)}(s_i^*, \mathbf{s}_{-i}) = c_i^{(\beta)}(\mathbf{s}^*) = r_1^{-C}$, so $x_i^{(\beta)}(s_i^*, \mathbf{s}_{-i}) - x_i^{(\beta)}(\mathbf{s}^*) = 0$. Accordingly, for the remainder of the proof we focus on the case $j > 1$.

Now, notice that conditioning on $C_i^{(\beta)}(\mathbf{s}) = C$ does not imply anything about i's ranking r, by independence of rankings. However, assuming $j > 1$ we know that r_1^{-C} was not available in profile $\mathbf{s}^*$, which could correlate with its value.

Consider then whoever got that candidate in profile $\mathbf{s}^*$ and let their ranking sample be f. The idea is that since she is also using a stochastically consistent technology, she will be "making the right choice" most of the time.

Now, we condition on an arbitrary pair of candidates k, ℓ, so that either $(r_1^{-C}, r_j^{-C}) = (k, \ell)$ or $(r_1^{-C}, r_j^{-C}) = (\ell, k)$. If we can show that for any such pair with $x(k) \geq x(\ell)$ the first event is the most likely, then we are done, since:

$$\mathbb{E}\left[x_i^{(\beta)}(s_i^*, \mathbf{s}_{-i}) - x_i^{(\beta)}(\mathbf{s}^*) \;\middle|\; \mathbf{x}, \beta, c_i^{(\beta)}(\mathbf{s}^*) \notin C_i^{(\beta)}(\mathbf{s})\right] =$$

$$1/2 \cdot \mathbb{E}_{C \sim C_i^{(\beta)}(\mathbf{s})} \sum_{j>1} \sum_{\substack{k,\ell \in [d] \\ k \neq \ell}} (x(k) - x(\ell)) \mathbb{P}\left(\mathcal{E}_{j,C}\right) \Big\{ \mathbb{P}\left((r_1^{-C}, r_j^{-C}) = (k, \ell) \mid \mathcal{E}_{j,C}\right)$$

$$- \mathbb{P}\left((r_1^{-C}, r_j^{-C}) = (\ell, k) \mid \mathcal{E}_{j,C}\right) \Big\}$$

Fix (condition) the arbitrary positions that k, ℓ could be ranked in rankings r and f, say u and $u+v$ in r and z, $z+w$ in f (note that they both have to "agree" on their relative positions by the conditioned event) and condition on everyone else's ranking (including the other positions of r and f) for simplicity of the argument.

To complete the proof we make use of the following inequality, for any $u < u+v, z < z+w \in [d]$ and for any $k \neq \ell \in [d]$ with $x(k) \geq x(\ell)$, when $\mathcal{E}'$ is an arbitrary event (consistent with the event $\mathcal{E}_{j,C}$) that may depend on the rankings of any other technology besides r and f, and on all positions of all candidates $c \notin \{k, \ell\}$ in rankings r and f.

$$\mathbb{P}\left(\begin{bmatrix} r_u \\ r_{u+v} \end{bmatrix} = \begin{bmatrix} k \\ \ell \end{bmatrix} = \begin{bmatrix} f_z \\ f_{z+w} \end{bmatrix} \;\middle|\; \mathcal{E}'\right) \geq \mathbb{P}\left(\begin{bmatrix} r_u \\ r_{u+v} \end{bmatrix} = \begin{bmatrix} \ell \\ k \end{bmatrix} = \begin{bmatrix} f_z \\ f_{z+w} \end{bmatrix} \;\middle|\; \mathcal{E}'\right) \tag{2}$$

The inequality is valid since distinct technologies are mutually independent and stochastically consistent. The independence implies that the left side factorizes as $\mathbb{P}\left(\begin{bmatrix} r_u \\ r_{u+v} \end{bmatrix} = \begin{bmatrix} k \\ \ell \end{bmatrix} \;\middle|\; \mathcal{E}'\right) \cdot \mathbb{P}\left(\begin{bmatrix} f_z \\ f_{z+w} \end{bmatrix} = \begin{bmatrix} k \\ \ell \end{bmatrix} \;\middle|\; \mathcal{E}'\right)$ and that the right side factorizes similarly. Stochastic consistency then implies that each factor on the left side is greater than or equal to the corresponding factor on the right side.

To see that inequality (2) suffices to complete the proof of the lemma, define $\mathcal{E}_{uvzw}^{k\ell}$ to be the event that $\{k, \ell\} = \{f_z, f_{z+w}\} = \{r_u, r_{u+v}\} = \{r_1^{-C}, r_j^{-C}\}$ (as *unordered* sets) and observe that $\mathcal{E}_{uvzw}^{k\ell}$ can be decomposed as a disjoint union of events $\mathcal{E}'$ of the type defined above. Hence, by taking a weighted average of inequalities (2) for all the constituent sets $\mathcal{E}'$ that compose $\mathcal{E}_{uvzw}^{k\ell}$, we obtain

$$\mathbb{P}\left(\begin{bmatrix} r_u \\ r_{u+v} \end{bmatrix} = \begin{bmatrix} k \\ \ell \end{bmatrix} = \begin{bmatrix} f_z \\ f_{z+w} \end{bmatrix} \;\middle|\; \mathcal{E}_{uvzw}^{k\ell}\right) \geq \mathbb{P}\left(\begin{bmatrix} r_u \\ r_{u+v} \end{bmatrix} = \begin{bmatrix} \ell \\ k \end{bmatrix} = \begin{bmatrix} f_z \\ f_{z+w} \end{bmatrix} \;\middle|\; \mathcal{E}_{uvzw}^{k\ell}\right).$$

Now we calculate

$$\mathbb{P}\left((r_1^{-C}, r_j^{-C}) = (k,\ell) \mid \mathcal{E}_{j,C}\right) - \mathbb{P}\left((r_1^{-C}, r_j^{-C}) = (\ell,k) \mid \mathcal{E}_{j,C}\right)$$
$$=$$
$$\sum_{\substack{u,v,z,w\in[d]\\ u+v>u\\ w+z>w}} \mathbb{P}\left(\mathcal{E}^{k\ell}_{uvzw} \mid \mathcal{E}_{j,C}\right) \left\{ \mathbb{P}\left(\begin{bmatrix} r_u \\ r_{u+v} \end{bmatrix} = \begin{bmatrix} k \\ \ell \end{bmatrix} = \begin{bmatrix} f_z \\ f_{z+w} \end{bmatrix} \;\middle|\; \mathcal{E}^{k\ell}_{uvzw} \right) \right.$$
$$\left. -\mathbb{P}\left(\begin{bmatrix} r_u \\ r_{u+v} \end{bmatrix} = \begin{bmatrix} \ell \\ k \end{bmatrix} = \begin{bmatrix} f_z \\ f_{z+w} \end{bmatrix} \;\middle|\; \mathcal{E}^{k\ell}_{uvzw} \right) \right\}$$

and deduce that the sum is non-negative because each summand is non-negative.

Finally, to conclude the proof we have to argue about "common" strategies, but observe that this proof can also be directly extended for that case too. There's only a few small changes, specifically in this case conditioning on $C_i^{(\beta)}(\mathbf{s}) = C$ might affect *some* of i's ranking, but crucially, stochastic consistency implies that it won't affect that it's more likely a pair specific pair of candidates in r^{-C} will be ranked correctly rather than not. Additionally, we won't be conditioning on everyone else's ranking, but only on those before i in both profiles, except the positions as defined previously. This again might affect some of i's ranking, but not the positions we care about (by the event), and so the arguments flow in the following fashion. If the firm who took candidate r_1^{-C} (in profile $\mathbf{s}$) is using another independent technology, the proof follows the same blueprint as previously. If the firm that took the candidate is using the same common technology, then the proof is simpler, as we just have to argue about the common technology's stochastic consistency.

□

Proposition 4. *Lemma 1 does not necessarily hold if firm i is using a stochastically consistent technology, but not her competitors.*

Proof. Assume the following for simplicity of the argument: $d = n + 3$, $\mathbf{x}$ is such that $x(i), i \in [d]$ is a decreasing sequence and β is such that i is last to choose ($\beta(i) = n$). Now, consider the following about the profiles $\mathbf{s}_{-i}$ and $\mathbf{s}^*_{-i}$: the former is such, that they always choose the best candidates, but never identify 1 to be such, meaning they will always hire the candidates of $[d] \setminus \{1, d-1, d\}$, leaving the rest for i. The latter, profile $\mathbf{s}^*_{-i}$, is not very well informed: they always choose the *worst* candidates, leaving $\{1, 2, 3\}$ for i to choose. Now, suppose i is employing a stochastically consistent ranking technology in $\mathbf{s}^*$, such as the Mallow's model with the Kendall-Tau distance. Her utility in profile $\mathbf{s}^*$ conditioned on the event of Lemma 1 is just $x(1)$. This is because the event is only true when she hires candidate 1, as candidates 2 and 3 are never available in $\mathbf{s}$. What is her conditional utility in profile $\mathbf{s}$ if she were to deviate to s_i^*? It must be less than $x(1)$, because there is a non zero probability she ranked d or $d-1$ above 1 (even though in profile $\mathbf{s}^*$ she would hire candidate 1). □

Acknowledgments. Erald Sinanaj was supported in part by AFOSR grant FA9550-23-1-0410. Éva Tardos was supported in part by AFOSR grants FA9550-23-1-0410 and FA9550-231-0068, and ONR MURI grant N000142412742.

References

1. Adamczyk, M., Sankowski, P., Zhang, Q.: Efficiency of truthful and symmetric mechanisms in one-sided matching. In: SAGT. pp. 13–24 (2014), DOI: 10.1007/978-3-662-44803-8_2
2. Baek, J., Makhdoumi, A.: The feedback loop of statistical discrimination. SSRN (2023), DOI: 10.2139/ssrn.4658797
3. Bhalgat, A., Chakrabarty, D., Khanna, S.: Social welfare in one-sided matching markets without money p. 87–98 (2011), DOI: 10.1007/978-3-642-22935-0_8
4. Bogomolnaia, A., Moulin, H.: A new solution to the random assignment problem. Journal of Economic Theory **100**(2), 295–328 (2001), DOI: 10.1006/jeth.2000.2710
5. Bommasani, R., Creel, K.A., Kumar, A., Jurafsky, D., Liang, P.S.: Picking on the same person: Does algorithmic monoculture lead to outcome homogenization? In: NeurIPS. vol. 35, pp. 3663–3678 (2022)
6. Christodoulou, G., Filos-Ratsikas, A., Frederiksen, S., Goldberg, P., Zhang, J., Zhang, J.: Social welfare in one-sided matching mechanisms. In: AAMAS. pp. 30–50 (May 2016), DOI: 10.1007/978-3-319-46882-2_3
7. E. Peajcariaac, Josip; L. Tong, Y.: Convex Functions, Partial Orderings, and Statistical Applications. Academic Press (Jun 1992)
8. Filos-Ratsikas, A., Frederiksen, S.K.S., Zhang, J.: Social welfare in one-sided matchings: Random priority and beyond. In: SAGT. pp. 1–12 (2014), DOI: 10.1007/978-3-662-44803-8_1
9. Jagadeesan, M., Jordan, M., Steinhardt, J., Haghtalab, N.: Improved bayes risk can yield reduced social welfare under competition. In: NeurIPS. vol. 36, pp. 66940–66952 (2023)
10. Jiang, S., Ndiaye, N., Vetta, A., Wu, E.: The price of anarchy of probabilistic serial in one-sided allocation problems. In: WINE. pp. 420–437 (2024), DOI: 10.1007/978-3-031-48974-7_24
11. Kendall, M.G.: A new measure of rank correlation. Biometrika **30**(1/2), 81–93 (1938), DOI: 10.2307/2332226
12. Kleinberg, J., Raghavan, M.: Algorithmic monoculture and social welfare. Proceedings of the National Academy of Sciences **118**(22) (2021), DOI: 1073/pnas.2018340118
13. Koutsoupias, E., Papadimitriou, C.: Worst-case equilibria. p. 404–413. STACS'99 (1999), DOI :10.1007/3-540-49116-3_38
14. Mallows, C.L.: Non-null ranking models. i. Biometrika **44**(1/2), 114–130 (1957), DOI: 10.2307/2333244
15. Peng, K., Garg, N.: Monoculture in matching markets. In: NeurIPS. vol. 37, pp. 81959–81991 (2024)
16. Roughgarden, T.: Intrinsic robustness of the price of anarchy. J. ACM **62**(5) (Nov 2015), DOI: 10.1145/2806883
17. Toups, C., Bommasani, R., Creel, K., Bana, S., Jurafsky, D., Liang, P.S.: Ecosystem-level analysis of deployed machine learning reveals homogeneous outcomes. In: NeurIPS. vol. 36, pp. 51178–51201 (2023)