Fall if it lifts your teammate: a novel type of candidate manipulation

Justin Kruger and Sebastian Schneckenburger

Abstract

We present a new interpretation of the traditional computational social choice framework, where what are traditionally the candidates are construed as the agents. The particular implementation in mind is the proposed system for determining the medal winners for sports climbing in the 2020 Olympic games. We consider the issues of ties and of potential manipulation with respect to this interpretation. Simulation results suggest that for the proposed system ties are unlikely to be a problem, but that there is at least potential for manipulation. We formalise the impossibility of manipulation with novel axioms. The strongest axioms lead to an impossibility along the lines of Arrow's impossibility, while a small weakening leads to a possibility.

1 Introduction

The 2020 Olympic Games in Tokyo will inaugurate ten new gold medals; one male and one female in each of five new events: karate, skateboarding, surfing, baseball and sports climbing. Of these, sports climbing did not exist in a unique competition format before its introduction as an Olympic event. Instead, there are three distinct types of competitive climbing: bouldering, lead-climbing and speed-climbing. Each discipline requires different skills and measures the performance of athletes using different methods. Thus sports climbing is to be a composite event, similar to the pentathlon. The novel event will have, however, its own novel system of determining the medal winners.

For Tokyo 2020 the International Federation of Sports Climbing (IFSC) has devised a combined format for sports climbing [11]. Twenty athletes—of which each country can have up to two representatives—will be involved in the main event. A qualification round reduces this number to six athletes who then compete in a final round to determine the medals. Both rounds proceed in the same manner: the contending athletes compete in all three disciplines, thereby producing three linear orders over the athletes. Each athlete is then assigned a score corresponding to the product of their rank in each discipline, where the rank of an athlete is the number of other athletes that defeat her plus one. These product scores provide a final ordering over the competitors, with lower scores being better. If two athletes receive the same score the tie is broken in favour of the athlete that performs better than the other in more disciplines.

As we have described it, the combined Olympic format may still lead to tied situations, because if more than two athletes receive the same product score the pairwise comparisons may form a cycle. Ties present a potential problem for both rounds of the competition. In the final round, it is desirable to have a single gold winner. In the qualifying round, tied situations may necessitate an extra method in order to determine which of the athletes progress—we provide an example to demonstrate the problem. The following table shows potential ranks of seven of the athletes after the qualification round, as well as bounds on the ranks of the other thirteen athletes given the specified results.

Athlete	Speed	Bouldering	Lead	Product
a	3	3	6	54
b	1	5	12	60
c	4	15	1	60
d	15	1	4	60
e	2	10	3	60
f	5	6	2	60
g	6	2	5	60
others	≥ 7	≥ 4	≥ 7	≥ 196

In this situation, *a* has the smallest product score and thus progresses to the final round. The unnamed athletes all have larger product scores than the seven named athletes, thus none of the unnamed athletes progress. There is a six way tie in the product scores of athletes *b* to *g*, thus one of these athletes must be eliminated from competing in the final round. For each pair of these athletes, we compare which athlete in the pair outperforms the other in more disciplines. This information can be concisely represented as a directed graph whose nodes are the tied athletes, with arrows to determine which athlete in each pair performs better. For instance, *b* performs better than *c* in speed and bouldering, thus we draw a directed arrow from *b* to *c*.

We display the graph representing the pairwise comparisons between the six tied athletes to the right. The particular arrangement of the athletes in a hexagon makes it clear that there is a cycle through the graph. Restricting attention to the six tied athletes, there is no *Condorcet loser*—no athlete that is pairwise defeated by all the others—as such it is not immediately obvious which of them should be eliminated.

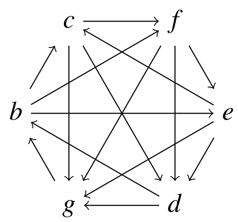

We continue developing our example to demonstrate another potential problem with the combined Olympic format: it may prompt deliberate bad performances. Let us suppose that, through some tie-breaking procedure, *b* is eliminated.[1] If the remaining athletes perform the same in the final round as they did in the qualification round, the ranks and products will be the following.

Athlete	Speed	Bouldering	Lead	Product
a	2	3	6	36
c	3	6	1	18
d	6	1	4	24
e	1	5	3	15
f	4	4	2	32
g	5	2	5	50

Suppose that *a* and *c* have the same nationality. According to the predicted results *a* will not win a medal, and certainly not the gold, while her teammate *c* is on course for the silver. However, if *a* deliberately performs worse than *c* in the speed competition and all other ranks remain the same, *c* will become the unique gold medal winner with a product score of 12.[2] National loyalty may lead *a* to *manipulate* in this manner, thereby spoiling *e*'s efforts which would otherwise have been enough for a gold medal.

Our example demonstrates two potential problems with the proposed competition format for climbing: firstly ties, and secondly a phenomenon that we refer to as manipulation. Resolving a

[1]The IFSC will use a "seeding list" to break ties that are not resolvable by pairwise comparisons. Such a seeding list is, in effect, an exogenous linear order tiebreaker. For the final round the ranking of the qualification round will be used as a seeding list; for the qualification round a seeding list based on the qualification system will be used [21].

[2]This would be feasible to perform in practice: in the speed competition in the final round, the third and fourth best speed-climbers will have at least one direct confrontation [11].

tie on factors other than the athletes' performances seems unfair, but may be necessary in order to determine which athletes progress to the final round. Ties are also undesirable in the final round—ideally, a gold medallist is unique. Concerning manipulation, an athlete who performs worse than she possibly may have done violates the Olympic spirit and cheapens the spectacle. Of course, only a very badly designed system would allow performing worse to lead to a better outcome for the manipulating athlete herself. We focus on the subtler case of *altruistic* manipulation, where the athlete aids a teammate by performing worse.

The rest of the paper is organised as follows. In Section 2 we define our framework, including two versions of altruistic manipulation. Section 3 is divided into two subsections in which we present our main results. Our simulation results in the first subsection show that, in fact, the probability of ties under the proposed method is quite low. On the other hand, it seems that there is a high potential for manipulation. This motivates a search for other methods that prevent the two kinds of manipulation; we discuss the theoretical possibility of such methods in the second subsection. Our theoretical results show that it is impossible to completely rule out manipulation, but we give a possibility for a plausible weakening. Section 4 provides a summary and further discussion; we discuss our framework's position with respect to other literature, possible extensions and other further work, and the relevance of our results to the climbing competition at the 2020 Olympics.

2 Definitions

Denote by $A = \{a, b, c, \dots\}$ the set of athletes and by $N = \{1, \dots, n\}$ the set of disciplines. We suppose that $m \geq 3$ and $n \geq 2$. Denote the set of linear orders on A by $\mathscr{L}$ and the set of total preorders on A by $\mathscr{W}$. We use $\succeq$ to denote a total preorder over the athletes, with $\succ$ the asymmetric part. All the athletes compete in each discipline $i \in N$, resulting in n linear orders $\succ_i$ over A. A *profile* that summarises the results for each discipline is denoted by $(\succ_1, \dots, \succ_n) = \boldsymbol{\succ} \in \mathscr{L}^N$. A *ranking function* f uses these results to produce a total preorder over the competitors: $f \colon \mathscr{L}^N \to \mathscr{W}$.

The ranking function should not be confused with the *rank* of competitors. Athletes have a rank for each discipline and for the output total preorder. Formally, for an ordering $\succeq$ over competitors, the rank of $a \in A$ is $r_{\succeq}(a) = |\{x \in A \colon x \succ a\}| + 1$. To simplify notation, for a discipline $i \in N$ we write $r_i = r_{\succeq_i}$. Because lower ranks are better, $\succ_i$ and the natural ordering on ranks are inverted: for all $x, y \in A$ and $i \in N$, $x \succ_i y$ iff $r_i(x) < r_i(y)$. The output can contain ties, though if two competitors are ranked first, no competitor is ranked second—a shared gold medal implies that no-one receives silver. We refer to athletes ranked first in the output as *winners*.

We can use a profile to make pairwise comparisons between athletes. Denote the distinct pairs of athletes by $D = \{(x, y) \mid x, y \in A, x \neq y\}$. Given a profile $\boldsymbol{\succ}$, define $\mathsf{ct}_{\boldsymbol{\succ}} \colon D \to \mathbb{Z}$ by

$$\mathsf{ct}_{\boldsymbol{\succ}}(x, y) = |\{i \in N \mid r_i(x) < r_i(y)\}| - |\{i \in N \mid r_i(y) < r_i(x)\}|.$$

For an arbitrary profile $\boldsymbol{\succ}$ the *weak majority relation* $T_{\boldsymbol{\succ}} \subseteq A \times A$ is defined by $x T_{\boldsymbol{\succ}} y$ iff $\mathsf{ct}_{\boldsymbol{\succ}}(x, y) \geq 0$. This relation is complete regardless of the parity of $|A|$. It may not be transitive: for a binary relation R we write R^+ for the transitive closure of R, the smallest transitive relation that contains R.

We now formally define the the proposed ranking function, insofar as it is determined by the profile of results. We call this function *inverse-Borda-Nash*,[3] denoted by $\mathsf{bn} \colon \mathscr{L}^N \to \mathscr{W}$. Define the

[3]This nomenclature is intended to be descriptive as it invokes Borda scores and the Nash product. The Nash product is sometimes described as providing a middle ground between the utility maximisation of additive methods and the maxi-min of egalitarian methods. However, this is not the case for the proposed method because the Borda scores are *inverted* i.e. smaller numbers are better. For addition, this inversion would have no effect, but this is not the case for multiplication. For example, according to inverse-Borda-Nash, an athlete with rankings (1,1,4) beats an athlete with (2,2,2); whereas for traditional Borda scores the opposite is true: (19,19,19) would be considered better than (20,20,17). It is seen as an advantage of the method that it favours specialists—it is preferred that the winner of the combined format is a potential winner of world-cups in some individual discipline, rather than a generalist [21]. We are not aware of any precedent for this method, this may be because it would become an "anti-fairness" approach when applied to social choice or social welfare.

binary relation $Q \subseteq A \times A$ by

$$xQy \quad \text{iff} \begin{cases} \prod_{i \in n} r_i(x) > \prod_{i \in n} r_i(y) \\ \text{or} \\ \prod_{i \in n} r_i(x) = \prod_{i \in n} r_i(y) \text{ and } xT_{\succ}y. \end{cases}$$

Define $\mathsf{bn}(\succ) = Q^+$; as Q is complete this is a total preorder.

2.1 Basic desiderata

There are some basic desiderata for a ranking function f. An athlete $x \in A$ clearly beats $y \in A$ in $\succ$ if for all $i \in N$, $x \succ_i y$. We say f satisfies the *clear winner condition* if whenever x clearly beats y in $\succ$, then for $\succeq = f(\succ)$ it is the case that $x \succ y$. We say f is *neutral* if permuting the competitors in the profile similarly permutes the competitors in the output ranking: for any permutation $\sigma: A \to A$, given $\succ$ and $\succ'$ such that for all $x, y \in A$, $i \in N$ $a \succ_i b \Leftrightarrow a \succ'_i b$, then $af(\succ)b \Leftrightarrow af(\succ')b$. Our last desiderata limits how much a single discipline can determine the winner. For a ranking function f, we say the gold is determined by $i \in N$ if, for any profile $\succ$ and writing $\succeq = f(\succ)$, $r_i(x) = 1$ implies $r_{\succeq}(x) = 1$. We say f is *non-determined* if the gold is not determined by any $i \in N$.

Each of these three desiderata can be linked to axioms from social choice theory. The clear winner condition is called, for e.g., the "Pareto criterion" [6, p. 42]. "Neutral" is standard terminology in social choice theory. A discipline that determines the gold may be thought of as a "weak top-dictator": an agent that can force his top ranked alternative to be among the top ranked alternatives in the output—note that this is a particularly weak version of a dictatorship.

2.2 Formal definitions of manipulation

The traditional definition of manipulation in social choice theory involves insincere representation of preference [8, 20]. Manipulation by an individual involves any possible change to one of the orderings in the profile. The situation we consider is slightly different: a potential manipulator has a position within each ordering and can manipulate, individually, by changing her position in multiple orderings. Also, she is restricted in the type of change she can make to each ordering; specifically, she can only make her own ranking worse: an athlete cannot perform better than her best.

The restriction on manipulation means that purely individual manipulation is not possible if the method satisfies a weak version of monotonicity. A ranking function f satisfies *ranking monotonicity* if for any two profiles $\succ, \succ' \in \mathscr{L}^N$ with $f(\succ) = \succeq$ and $f(\succ') = \succeq'$, for any competitor $a \in A$, if for all $i \in N$, $x \in A \backslash \{a\}$ and $y \in A$ it is the case that $x \succ_i y$ implies $x \succ'_i y$, then $r_{\succeq}(a) \leq r_{\succeq'}(a)$. In words, this requires that an individual cannot increase their own output ranking by performing worse.

Inverse-Borda-Nash satisfies ranking monotonicity, thus prevents purely individual manipulation. However, two athletes of the same nationality can compete in the sports climbing event: an athlete may be able to altruistically manipulate for their teammate. In particular, a manipulator may be able to improve their teammate's ranking without worsening their own output ranking.

Definition 1 (Non-sacrificial manipulation)**.** *Let $f(\succ) = \succeq$ and $f(\succ') = \succeq'$, and $a, b \in A$. Athlete a can manipulate without sacrifice, for athlete b, from the profile $\succ$ to the profile $\succ'$ if*

1. *for all $i \in N$, $x \in A \backslash \{a\}$ and $y \in A$, $x \succ_i y$ implies $x \succ'_i y$*
2. *$r_{\succeq'}(b) < r_{\succeq}(b)$*
3. *$r_{\succeq'}(a) \leq r_{\succeq}(a)$.*

Such a manipulation is strictly *without sacrifice if it also satisfies*

4. *$|\{x \in A : x \succeq' a\}| \geq |\{x \in A : x \succeq a\}|$.*

A ranking function f prevents manipulation without sacrifice if there is no pair of profiles and pair of competitors that can manipulate without sacrifice. The strictness condition makes it harder to manipulate: the idea is that an athlete prefers to be uniquely ranked in a position than to share this ranking with multiple athletes. This makes preventing strictly non-sacrificial manipulation a weaker condition than preventing all kinds of non-sacrificial manipulation. This will make the difference between an impossibility and possibility result.

The example of manipulation in the introduction is not without sacrifice. Instead, the idea is that the manipulating athlete recognises that she cannot get a higher ranking than her teammate, but can nonetheless manipulate to aid her teammate. We call this "spoiler" manipulation, this refers to the fact that a poorly ranked athlete spoils the fair result concerning other, better ranked, athletes.

Definition 2 (Spoiler manipulation)**.** *Let $f(\succ) = \succeq$ and $f(\succ') = \succeq'$, and $a, b \in A$. Athlete a can spoil, for athlete b, from the profile $\succ$ to the profile $\succ'$ if*

1. *for all $i \in N$, $x \in A \backslash \{a\}$ and $y \in A$, $x \succ_i y$ implies $x \succ'_i y$*
2. $r_{\succeq'}(b) < r_{\succeq}(b)$
3. $r_{\succeq'}(b) < r_{\succeq}(a)$.

Where it is not necessary to specify b or the profiles, we say that a spoils.

Ranking monotonicity prevents the special case of manipulation where an athlete manipulates for herself. That is, if f prevents manipulation (strictly) without sacrifice, or if f prevents spoiler manipulation, then f satisfies ranking monotonicity.

3 Results

Inverse-Borda-Nash satisfies the three basic desiderata.

Proposition 1. *Inverse-Borda-Nash satisfies the clear winner condition, is neutral, and is non-determined.*

Proof. Clear winner: if an athlete is ranked better than another in all disciplines, it must have a smaller product of ranks, therefore will be ranked better in the output.

Neutrality: if we permute athletes, we also permute their product scores and the relation $T_{\succ}$.

Non-determined: for an arbitrary discipline $i \in N$, take a profile where some athlete $a \in A$ comes first in all other disciplines and second in this discipline: $r_i(a) = 2$ and $r_j(a) = 1$ for all $j \in N$, $j \neq i$. If $n > 2$, then a is the unique winner, thus the discipline does not determine the gold. For $n = 2$, the following profile shows that i does not determine the gold (we use the assumption that $m \geq 3$):

Athlete	i	$j \neq i$	Product
a	2	1	2
b	1	3	3
c	3	2	6
others	≥ 4	≥ 4	≥ 16

□

However, inverse-Borda-Nash may produce ties and is potentially manipulable. We determine whether these issues are likely to occur by simulating competitions in the next subsection. We apply a theoretical analysis in the subsection that follows.

Culture	Ties	Spoiler manipulation	Without sacrifice	Strict without sacrifice	Any manipulation
Impartial	632	37,730	47,807	47,326	59,660
Positive cor.	779	13,792	43,723	41,597	46,964
Negative cor.	526	44,826	48,741	48,350	63,151

Table 1: The number of randomly generated profiles that involved ties, were subject to spoiler manipulation, were subject to manipulation without sacrifice, were subject to manipulation strictly without sacrifice, and that were subject to any of the manipulations that we define. 100,000 profiles were generated for each culture.

3.1 Simulations

Our simulations suggest that for inverse-Borda-Nash, although ties are unlikely to be a problem, potential for manipulation occurs with a high probability.

We generated profiles with six athletes and three disciplines, the same numbers as in the final round of the Olympics sports climbing competition. The generated profiles form three groups: in the first group, for each discipline every possible linear order is equally likely—this is the *impartial culture* [22]. For profiles in the second group there is a positive correlation in an athlete's results across the three disciplines. In the final group there is positive correlation between two disciplines and negative correlation with the third. This third culture conforms best to our actual expectations for the competition because the two disciplines of bouldering and lead climbing have an intersection of athletes at the top level, whereas top level speed climbers do not typically compete in the other disciplines.

For an impartial culture profile we independently select each of the linear orders uniformly at random from the set of all possible linear orders. Concerning the positively correlated culture, we also select the linear orders independently from each other, in the following manner. Label the athletes from a to f. Each linear order is constructed by successively adding athletes to partially completed linear orders, starting with the linear order consisting of only athlete a. We then flip a fair coin to decide whether b is placed before or after a, thereby creating a linear order over the two athletes. Next we add c: with probability $1/2$ we place her last in the partially constructed linear order; otherwise, we flip a second fair coin to see whether she is placed second in the linear order or first. The same idea is then applied to the other athletes: for each we successively flip a fair coin to determine whether or not the athlete is placed at an incrementing position in the partially constructed linear order, starting from the last position. A negatively correlated profile is created by taking a positively correlated profile and reversing the linear order of the last discipline.

We randomly generated 100,000 profiles of each type. A profile counts as tied if at least one tie occurs at any ranking level—we do not count the number of distinct ties nor how many athletes are involved in each tie. To count manipulations, we first randomly pair the athletes into three disjoint pairs. A profile counts as manipulable if at least one of the pairs can manipulate. We perform the count separately for spoiler manipulation, manipulation without sacrifice, manipulation strictly without sacrifice, and for any type of manipulation. The results are presented in Table 1.

According to our models, it is very unlikely that there will be a tie at any level of the output total preorder in the final round of the competition. We also ran simulations for twenty athlete profiles obtaining similar results.[4] This strongly supports the idea that a tie in the actual competition is very unlikely to occur: note that each of our models exhibits a high degree of symmetry; one would expect that such symmetries would be the most likely to cause tied situations. Indeed, it has been shown

[4]For profiles with twenty athletes, of the 100,000 profiles we generated for each culture, 208 profiles had ties for the impartial culture, 1108 profiles had ties for the positively correlated culture, and only 92 profiles had ties for the negatively correlated culture that conforms best to our expectations for the actual competition.

that the impartial culture maximises the probability for majority cycles [22], one of the necessary conditions for a tie. However, from our simulations we see more ties for the positively correlated culture: this is perhaps because two opposing criteria need to be fulfilled for there to be a tie; there need to be majority cycles, but these must occur among athletes with the same *scores*. Regardless, the incidence of ties is low even for the positive culture. Of the three cultures, we see fewest ties ties in the negatively correlated culture which best represents our expectations for the competition.

In contrast to the low incidence of ties, there does seem to be a high potential for manipulation, of both kinds. For each culture approximately half the profiles are manipulable.[5] We make two further observations: first, the incidence of *non*-strict manipulation without sacrifice is very small—the value obtained when subtracting the value of column five from the value of column four. A loose interpretation is that for inverse-Borda-Nash there is not much difference between the stronger and weaker versions of the without sacrifice axiom. Second, spoiler manipulation seems less likely under the positively correlated culture. An intuitive explanation for the lower incidence of spoiler manipulation for the positively correlated culture is the following: for this culture it is more likely that one athlete in a pair will *always* be ranked above their teammate, in which case the lower ranked athlete cannot spoil. Nevertheless, even for the positive culture there is a non-negligible potential for spoiler manipulation (more than 10% of the generated profiles).

3.2 Theoretical results

We want to define a method that satisfies our desiderata and completely prevents both forms of manipulation. Unfortunately, it is impossible to completely succeed in this task.

Theorem 1. *No ranking function prevents spoiler manipulation, prevents manipulation without sacrifice, satisfies the clear winner condition, and is non-determined.*

Proof. We show something slightly stronger: if both types of manipulation are prevented and the clear winner condition is satisfied then there is a *top-dictator*: a discipline $i \in N$ such that, for $a \in A$ such that $r_i(a) = 1$, $a \succ x$ for all $x \neq a$. This directly implies that the function is determined. The proof follows the structure of the proof of Arrow's theorem by Reny [17].

Take an arbitrary ranking function f that prevents spoiler manipulation and manipulation without sacrifice and that satisfies the clear winner condition. Consider any profile where a comes first in all disciplines and b comes last. By the clear winner condition a must be ranked first and b last. We express this fact as in the diagram to the right.

N		
a	$\mapsto$	a
$\mid$		$\mid$
b		b

Now consider moving b up in the first discipline. So long as b does not cross above a, a must still be uniquely ranked first, as otherwise the agent c that b becomes ranked above can spoil for the new winner from P_1 to P_2.

P_1

1	>1		
a	a	$\mapsto$	a
$\mid$	$\mid$		$\mid$
c	$\mid$		
b	$\mid$		
$\mid$	$\mid$		
$\mid$	b		

P_2

1	>1		
a	a	$\mapsto$	a
$\mid$	$\mid$		$\mid$
b	$\mid$		
c	$\mid$		
$\mid$	$\mid$		
$\mid$	b		

[5] We also tested profiles with twenty athletes for manipulation. Each culture resulted in higher counts of potential manipulation than in the six athlete case. Of course, to fully address the issue of manipulation in the qualification round would require other modifications: here a manipulation is only desirable if it moves the target agent below the sixth place threshold; more fundamentally we require an argument for why the the non-manipulated profile is common knowledge.

If we rank b above a, either a remains the unique winner or there is some other set of winners. If we continue to rank b first successively for the remaining disciplines, eventually b becomes the unique winner by the clear winner condition—in particular when b is ranked first in all disciplines—thus the second disjunct of the previous sentence happens at some point; there is a profile that outputs a top, while if b is moved above a in one discipline the set of winners is X with some $x \in X$ such that $x \neq a$. Label the discipline for which this happens i^* and label the respective profiles as P_3 and P_4.

P_3

$< i^*$	i^*	$> i^*$
b	a	a
a	b	$\vert$
$\vert$	$\vert$	$\vert$
$\vert$	$\vert$	b

$\mapsto$ a, $\vert$

P_4

$< i^*$	i^*	$> i^*$
b	b	a
a	a	$\vert$
$\vert$	$\vert$	$\vert$
$\vert$	$\vert$	b

$\mapsto$ b, $\vert$

We know that $a \notin X$, otherwise a could spoil without sacrifice for x from P_3 to P_4. This implies that $b \in X$, as otherwise b could spoil for a from P_4 to P_3. This implies that $x \notin X$ for $x \neq a, b$, as otherwise b could manipulate without sacrifice for x from the profile where b is the unique winner.

In P_4 we can move a down in the profile without changing the output winner b (otherwise a could spoil), we display this as P_5. Create P_6 from P_5 by moving a up one place in discipline i^*.

P_5

$< i^*$	i^*	$> i^*$
b	b	$\vert$
$\vert$	a	$\vert$
$\vert$	$\vert$	$\vert$
$\vert$	$\vert$	a
a	$\vert$	b

$\mapsto$ b, $\vert$

P_6

$< i^*$	i^*	$> i^*$
b	a	$\vert$
$\vert$	b	$\vert$
$\vert$	$\vert$	$\vert$
$\vert$	$\vert$	a
a	$\vert$	b

$\mapsto$ a, $\vert$

Athlete a must the unique winner in P_6. First, note that if neither a nor b were ranked first for P_6, then a can spoil for b from P_6 to P_5. If b is ranked first but not uniquely ranked first, then b can spoil without sacrifice from P_5 to P_6. If b is uniquely ranked first, then at some point in stepwise changes from P_6 to P_3 some other athlete must perform a spoiler manipulation. Thus as b is not ranked first a is amongst the winners. If a were not unique, a could spoil without sacrifice from P_3 to P_6.

Take some third alternative $c \neq a, b$. The profile P_7 is obtained from P_6 by moving b and c down in the profile. Here the unique winner is still a, as otherwise b or c could spoil. Create P_8 by moving a to be ranked last in all disciplines except i^*.

P_7

$< i^*$	i^*	$> i^*$
$\vert$	a	$\vert$
$\vert$	$\vert$	$\vert$
c	$\vert$	c
b	c	a
a	b	b

$\mapsto$ a, $\vert$

P_8

$< i^*$	i^*	$> i^*$
$\vert$	a	$\vert$
$\vert$	$\vert$	$\vert$
c	$\vert$	c
b	c	b
a	b	a

$\mapsto$ a, $\vert$

In the profile P_8, alternative c is a clear winner over b, so b cannot be ranked first. If a were not ranked first then b could spoil for a from P_8 to P_7. If any other athlete is ranked first, then a can manipulate without sacrifice from P_7 to P_8. Thus a must be the unique winner in P_8.

In general, for any profile where a wins in discipline i^*, a must be uniquely ranked first in the output, as otherwise there would be some chain of changes from P_8 to the profile in question, one of which would be a spoiler manipulation for the new winning athlete. As a is arbitrary, for each alternative x there is a discipline i_x such that whenever x wins in i_x, x is uniquely ranked first. As two alternatives x and y cannot both be ranked first, $i_x = i_y$ for all $x, y \in A$, thus i^* is a top dictator. □

The proof closely follows Reny [17], who presents Arrow's impossibility and the Gibbard-Satterthwaite theorem side by side. Although we consider manipulation, the result is, in terms of its formal shape, closer to Arrow's result than to the Gibbard-Satterthwaite result. Requiring the impossibility of both forms of manipulation replaces the axiom of independence of irrelevant alternatives (IIA), though this requirement does not *imply* IIA. Consider the ranking function that always returns only two ranks, one consisting of only the alternative a and the other consisting of all the other alternatives, such that a is ranked first iff a is first in all the disciplines; this violates IIA but prevents both kinds of manipulation. Alongside the fact that non-determined is a weakening of non-dictatoriality, this means that our impossibility is not simply a corollary of Arrow's theorem.

We cannot satisfy all our desiderata simultaneously. However, if we weaken manipulation without sacrifice to manipulation strictly without sacrifice there are methods that work. The method we define proceeds in stages, determining the top ranked candidates then removing them from the profile. It may be thought of as a back-to-front version of *instant runoff voting* [26, p. 37] applied using a *majority quota rule*. Also cf. the *Coombs rule* [10]. If an athlete is ranked first in strictly more than half the disciplines, then she is the unique winner with respect to the athletes in the profile. Otherwise, any athlete that has at least one first place ranking in the profile is a joint winner. The winners are removed from the profile, and the procedure repeats. We name this *iterative first place elimination*, $\mathsf{ifpe} : \mathscr{L}^N \to \mathscr{W}$. Formally, for an arbitrary profile $\succ$, let

$$\mathsf{win}(\succ) = \begin{cases} \{a\} & \text{if } \exists a \in A,\ |\{i \in N : r_i(a) = 1\}| > n/2, \\ \{x \in A : \exists i \in N, r_i(x) = 1\} & \text{otherwise.} \end{cases}$$

Let $\succ^1 = \succ$. For $t \geq 1$, recursively define $\succ^{t+1}$ as the restriction of $\succ^t$ to $A \backslash \mathsf{win}(\succ^t)$. Writing $\succeq = \mathsf{ifpe}(\succ)$, for $x, y \in A$, define $x \succeq y$ iff there are integers s, t such that $s \leq t$ and $x \in \mathsf{win}(\succ^s)$ and $y \in \mathsf{win}(\succ^t)$.

Proposition 2. *Iterative first place elimination prevents spoiler manipulation, prevents manipulation strictly without sacrifice, satisfies the clear winner condition, satisfies neutrality, and, for $n \geq 3$, is non-determined.*

Proof. Prevents spoiler manipulation: an athlete cannot affect any of the partial profiles $\succ^t$ starting from $t = 1$ until the profile where she is ranked first in one of the disciplines. Consider the partial profiles for which she is ranked first. There are two possibilities. (1) The athlete is not a winner for the partial profile, thus a different athlete is ranked first in more than half the disciplines; this other athlete will be the winner no matter how the putative manipulator changes her ranking. (2) The athlete is a winner for the partial profile, thus she cannot spoil because she does as well as the remaining athletes.

Prevents manipulation strictly without sacrifice: an athlete a cannot affect athletes that get better output ranks. Let $\succ$ be the partial profile for which $a \in \mathsf{win}(\succ)$. First suppose a is ranked first in more than half the disciplines: if she performs worse in enough of these disciplines she will no longer be the unique winner, but such a manipulation is not strict. Otherwise, she will be removed from the profile in the next step, thus any changes to her ranking do not affect the output. Second suppose a is ranked first in less than half the disciplines. If a performs worse in a discipline i for which $r_i(a) > 1$, this will not affect the output ranking as a is removed from the profile in the next round. If a performs worse in a discipline i for which $r_i(a) = 1$, there are three possibilities. (1) A different athlete becomes the unique winner, thus a is ranked lower in the output. (2) A new athlete becomes a winner, in which case the manipulation is not strict. (3) The winners remain the same, thus the same athletes will be removed from this profile and the output will not change.

Clear winner: if a is ranked better than b in all disciplines, it is not possible that b is ranked first in a partial profile while a is still contained in the profile.

Neutrality: permuting the athletes in the profile will result in permuted sets $\mathsf{win}(\succ)$.

Non-determined: here we require the condition that $n \geq 3$; for arbitrary $i \in N$ consider the profile

Athlete	$\{i\}$	$N \backslash \{i\}$
a	2	1
b	1	2
others	≥ 3	≥ 3

□

Iterative first place elimination is unsatisfactory because it is *indecisive*, where we use the (slightly imprecise) term "decisiveness" to refer to a measure of how often ties are produced in the output. At the sharpest end, a maximally decisive method would always produce a linear order. Requiring this level of decisiveness recreates the impossibility because it makes manipulation without sacrifice equivalent to the strict manipulation without sacrifice. However, maximal decisiveness is arguably too strong a condition: for completely symmetric profiles, it seems reasonable in practice that conditions external to the profile break the ties. We are not aware of an axiomatic analysis of decisiveness that provides suitably weaker definitions that further investigation may be based upon.

4 Final remarks

In this paper we propose a novel interpretation of Arrow's traditional social choice framework involving the aggregation of linear orders. Under this interpretation what are traditionally thought of as candidates are the agents of the model. These agents can strategize in a specific manner: they can worsen their own position within one or more of the input linear orders. This interpretation captures the problem of aggregating multiple ranked competitions. In particular we consider the method proposed for determing the medal winners for climbing at the 2020 Olympics. Simulations suggest that, although ties are unlikely to occur, this method is potentially open to manipulation. Although it is impossible to completely rule out the least restrictive definitions of manipulation, a small assumption about how athletes are willing to manipulate means that non-manipulable methods are possible. The method that demonstrates this possibility is, however, very susceptible to ties.

Our interpretation is novel to the best of our knowledge. Other work concerning manipulation in sports competitions includes work concerning manipulating seedings [18], and tends to be of a more operations-research nature than social-choice-theoretic, see [23] for a survey. We are not aware of other work that explicitly considers candidates as agents in the way that we do—our work is distinct from the strand of literature which considers manipulation by strategic candidacy [5, 7, 14]. Of course, there are similarities between our results and more traditional work in social choice theory, and there may be implicit connections that we have missed, for example with the definitions of *Condorcet independence of irrelevant alternatives* [25] or *one-way monotonicity* [19].

Our interpretation fits well into Arrow's framework. Arguably, the problem of aggregating multiple disciplines is better served by this framework than typical problems of social choice theory. The linear order profile is the input in practice. There are no questions, as there are for social choice theory, about whether eliciting full linear orders is problematic, let alone whether linear order preferences are suitable or even sensible—cf. competing approaches like approval voting [15] and majority judgment [1].[6] The required output is also obviously a total preorder, whereas in so-

[6] Rather than linear order profiles, it would be possible to use a method that assigns points based upon individual performances. Thus the final score for an athlete would be necessarily independent of the performance of other athletes, sidestepping our issue of manipulation. Such a method is used for the modern pentathlon. This approach was discarded by the IFSC because (1) it is too complex for spectators and (2) it is difficult, perhaps impossible, to assign points in a balanced way across the disciplines [21]. We thus take it is as given that the input is ordinal.

Interestingly, Balinski and Laraki [1] use the example of Olympic figure skating as part of their argument against the ordinal approach in social choice. In the past the ranking of skaters was produced by aggregating multiple *ordinal* rankings given by multiple judges. The particular method has since been replaced, and it is argued that this is because it violates cial choice theory often what is desired is a *choice*, requiring a "social choice function" as opposed to a "social welfare function". Sometimes it is not obvious that manipulation is actually undesirable for social choice theory, especially when one considers iterative manipulation [16]. Indeed, "manipulation" is a misnomer, a better term would be *strategic behaviour*. In contrast, for sports competitions manipulation is aptly named and clearly undesirable in and of itself, whether because it goes against the spirit of the competition or because it cheapens the spectacle. Concerning the information requirements for manipulation, we have argued that the qualification round can be used as a proxy for the results in the final round. This is unrealistic—the athletes will not perform exactly the same—however it is certainly not less realistic than the traditional Gibbard-Satterthwaite assumption of common knowledge of all preferences of all agents.

There is one way in which our interpretation has a slightly different focus from that of traditional social choice: it stresses the importance of having a minimal *rank*; an athlete is only concerned with the number of athletes ranked *strictly higher* than her in the output total preorder. Another difference is a particular importance on the "decisiveness" of the ranking method; how often ties are output at any ranking level. Authors often sidestep the issue of ties in order to obtain their main results, by supposing that there is an exogenous linear order tiebreaker or by restricting the output to linear orders [26, p. 33], but this is obviously unsatisfactory for our purposes because it is the issue of ties itself that we are interested in. Alternatives approaches such as using a randomised mechanism to break ties or dealing directly with set-valued outcomes [2] are similarly unsatisfactory. We are not aware of a good reference for this subject.

Our impossibility result is perhaps not very exciting. Even though it is a *non-trivial* adaptation of Arrow's famous impossibility, because preventing manipulation is not equivalent to independence of irrelevant alternatives, it remains only an adaptation. For completeness, we note that the impossibility is tight for the four conditions. *Dictatorships*, where the ranking of a single discipline are copied, violate only non-determination. *Constant functions* violate only the clear winner condition, except the function that always ranks every athlete first. (Constant functions also violate neutrality, but this is not included in the impossibility.) Our method of *iterative-first-place-elimination* only allows manipulation without sacrifice. Finally, we sketch an upside-down variant of instant runoff voting that only allows spoiler manipulation: at stage t, remove the athlete who is ranked last in discipline t modulo n, and rank this athlete below the other athletes remaining in the profile. A formal definition and proof of properties is in the appendix.

Our positive result is interesting in part because of its unsuitability: it would certainly produce too many ties to be useful in practice. On the other hand, completely prohibiting ties recreates the impossibility. It would be interesting to determine if there is a satisfactory middle ground, but as we have already noted there seems to be an intriguing gap in, at least our knowledge of, the literature.

One of the originating ideas of *computational* social choice is that, even if manipulation is possible in theory, it might be computational hard to determine a strategy for manipulation, thus manipulation is unfeasible in practice [3]. Of course, it is obviously easy in practice to find manipulations for the proposed competition format with its fixed parameters of three disciplines and six or twenty athletes, otherwise our simulations would still be running. The complexity of finding a worthwhile manipulation for a pair—spoiler or without sacrifice—is polynomial in the number of athletes, one need only check the result for the (less than) 2^n profiles where the manipulator does just worse than her teammate for each possible subset of the disciplines. We do not know whether or not varying the amount of disciplines leads to a hardness result.

An interesting extension of our model would be to apply the *protocol approach*, where one considers partial revelation of the profile in a sequential manner. This is precisely how the Olympic sports climbing event will unfold, though it should be noted that there will be measures put in place to isolate the athletes from the partial results—perhaps a determined manipulator will find a way around these. A sequential extension would also be applicable to other competition formats. There

independence. Balinski and Laraki thus go in the opposite direction to us: they use experience from the Olympics and apply it to social choice theory.

is already a literature of related results concerning necessary and possible winners (stemming from [13]) which should greatly aid the development of such an extension.

One can postulate a variant of Murphy's law for sport competitions: "If a sport competition is susceptible to a certain kind of manipulation, some day some athlete will manipulate in that way". There is evidence for this law; cheating occurs in the Olympics, and specifically deliberate bad performances [12, 24]. The combined format for sports climbing is susceptible to a particular kind of altruistic manipulation that has been observed to occur in Formula 1 [9].[7] So do we expect a wave of manipulation at Tokyo 2020? Probably not: in favour of the proposed method for climbing, we do not believe that such manipulation is likely to happen in practice. Athletes will be held in isolation during several rounds, so unless there is some information leak they won't have enough information to successfully manipulate according to a protocol approach. This will not prevent athletes from using the results of the qualification round as a proxy, as we suggest in the example in the introduction. However, in practice this proxy will not be reliable. Lead-climbing routes and boulder problems are not standardised. For instance some problems can be easier for tall athletes, while others are easier for small athletes; the problems in the qualifying round may favour one type and those in the final round the other. Similarly, an athlete may rise to the challenge and perform better in the final round than in the qualifying round, or vice-versa crack under the pressure. Given the unpredictable nature of the final event, it will be hard to ensure that the intended manipulation does not have unintended effects.

A Method that only allows spoiler manipulation

The following method is closely related to the Combs rule [10]. Define *iterative successive last removal*, $\mathsf{isr}\colon \mathscr{L}^N \to \mathscr{W}$ as follows. For an arbitrary profile $\succ$, let $\mathsf{lose}_t(\succ) = \{a \in A : r_t(a) \text{ is maximal}\}$. Let $\succ^1 = \succ$, and for $t \geq 1$ recursively define $\succ^{t+1}$ as the restriction of $\succ^t$ to $A \backslash \mathsf{lose}_{t \bmod n}(\succ^t)$. Writing $\succ = \mathsf{isr}(\succ)$, for $x, y \in A$, define $x \succeq y$ iff there are integers $s, t \leq m$ such that $s \geq t$ and $x \in \mathsf{lose}_{s \bmod n}(\succ^s)$ and $y \in \mathsf{lose}_{t \bmod n}(\succ^t)$.

Proposition 3. *Iterative successive last removal prevents manipulation without sacrifice, satisfies the clear winner condition, satisfies neutrality and is non-determined.*

Proof. Prevents manipulation without sacrifice: suppose an athlete "manipulates" by performing worse in a profile but also that she does not get a worse output ranking. Thus she is removed at the same point t and has output rank $m - t + 1$. All the partial profiles after this point will be the same as in the non-manipulated case. As she was not removed before t, this means that for all the partial profiles at stage $s < t$ she was not ranked last in discipline s modulo n, this means that she did not change the athlete who was ranked last in this discipline, thus the loser at this stage will be the same.

Clear winner: if a is better than b in all disciplines then a cannot be removed before b.

Neutrality: permuting the athletes in the profile results in permuted sets $\mathsf{lose}(\succ)$.

Non-determined: for disciplines $i \neq 1$, consider the profile where the athlete ranked first in i is ranked last in 1. For discipline 1 consider the profile where the athlete ranked first in 1 is ranked last in 2. □

[7] As noted by a reviewer, for Formula 1, the type of manipulation we focus on in this paper may be considered to be less "important" than the phenomenon where a racer blocks the passage of cars of rival teams in order to preserve the advantage for his teammate. This perhaps exemplifies a distinction between manipulation and strategic behaviour. A single "blocking move" is sanctioned in Formula 1 [4]. Allowing blocking may be argued to be desirable, as it adds a strategic level for the competitors and increases tension for the spectators: if one agrees with this argument one may consider blocking as strategic behaviour rather than manipulation. The type of manipulation we focus on, though perhaps rare, also occurs in Formula 1 [9]. However our manipulation is more clearly undesirable—the reviewer used the term "scandalous".

Teamwork also plays an important role in other seemingly individualistic competitions. For example, in various cycling events, teammates draft behind each other. These strategic considerations are domain specific and would be difficult to treat in an abstract manner, as we do with our definition of manipulation in this paper. We do not see how the tools of computational social choice could be used for such specific cases.

References

[1] Michel Balinski and Rida Laraki. *Majority Judgment: Measuring, Ranking, and Electing*. MIT press, 2011.

[2] Salvador Barberà, Walter Bossert, and Prasanta K Pattanaik. Ranking sets of objects. In *Handbook of utility theory*, pages 893–977. Springer, 2004.

[3] John J. Bartholdi, Craig A. Tovey, and Michael A. Trick. The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6(3):227–241, 1989.

[4] Andrew Benson. Formula 1 bosses clarify blocking rule. *BBC*, 4 January 2012. `https://www.bbc.com/sport/formula1/16416932` (accessed 12 April 2018).

[5] Markus Brill and Vincent Conitzer. Strategic voting and strategic candidacy. In *Proceedings of the 29th conference on Artificial Intelligence (AAAI)*, pages 819–826, 2015.

[6] Donald E. Campbell and Jerry S. Kelly. Impossibility theorems in the arrovian framework. In Kenneth J. Arrow, Amartya K. Sen, and Kotaro Suzumura, editors, *Handbook of Social Choice and Welfare*, volume 1, chapter 1, pages 35–94. Elsevier, 2002.

[7] Bhaskar Dutta, Matthew O. Jackson, and Michel Le Breton. Strategic candidacy and voting procedures. *Econometrica*, 69(4):1013–1037, 2001.

[8] Allan Gibbard. Manipulation of voting schemes: a general result. *Econometrica*, 41(4):587–601, 1973.

[9] Gpupdate. Barrichello admits Ferrari team orders. *gpgupdate.net*, 25 November 2008. `https://www.gpupdate.net/en/f1-news/203773/barrichello-admits-ferrari-team-orders/` (accessed 28 February 2018).

[10] Bernard Grofman and Scott L. Feld. If you like the alternative vote (aka the instant runoff), then you ought to know about the Coombs rule. *Electoral studies*, 23(4):641–659, 2004.

[11] International Federation of Sports Climbing. Combined 2017 format for IFSC events and olympic games. *ifsc-climbing.org*, 2017. `https://www.ifsc-climbing.org/images/World_competitions/Event_regulations/170314_Combined-Olympic-Format.pdf` (accessed 28 February 2018).

[12] Tony Jimenez. Cycling - I crashed on purpose, says victorious Briton. *Reuters*, 3. August 2012. `https://uk.reuters.com/article/uk-oly-cycl-ctmspr-hindes-day7/cycling-i-crashed-on-purpose-says-victorious-briton-idUKBRE87201X20120803` (accessed 28 February 2018).

[13] Kathrin Konczak and Jérôme Lang. Voting procedures with incomplete preferences. In *Proceedings of the IJCAI-05 Multidisciplinary Workshop on Advances in Preference Handling*, 2005.

[14] Jérome Lang, Nicolas Maudet, and Maria Polukarov. New results on equilibria in strategic candidacy. In *Proceedings of the 6th International Symposium on Algorithmic Game Theory*, pages 13–25, 2013.

[15] Jean-François Laslier and M. Remzi Sanver. *Handbook on approval voting*. Springer Science & Business Media, 2010.

[16] Omer Lev and Jeffrey S. Rosenschein. Convergence of iterative voting. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*, pages 611–618, 2012.

[17] Philip J. Reny. Arrow's theorem and the Gibbard-Satterthwaite theorem: a unified approach. *Economics Letters*, 70(1):99–105, 2001.

[18] Tyrel Russell and Peter Van Beek. An empirical study of seeding manipulations and their prevention. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence*, pages 350–356, 2011.

[19] M. Remzi Sanver and William S. Zwicker. One-way monotonicity as a form of strategy-proofness. *International Journal of Game Theory*, 38(4):553–574, 2009.

[20] Mark A. Satterthwaite. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10(2):187–217, 1975.

[21] Jérôme Meyer (Sport Director IFSC). Interviewed by Sebastian Schneckenburger. Personal communication, 2018.

[22] Ilia Tsetlin, Michel Regenwetter, and Bernard Grofman. The impartial culture maximizes the probability of majority cycles. *Social Choice and Welfare*, 21(3):387–398, 2003.

[23] Mike Wright. OR analysis of sporting rules–a survey. *European Journal of Operational Research*, 232(1):1–8, 2014.

[24] Staff writers. Olympic badminton players charged with trying to lose their games. *The Guardian*, 31 July 2012. `https://www.theguardian.com/sport/2012/aug/01/olympic-badminton-players-charged-lose` (accessed 28 February 2018).

[25] Ning N. Yu. A quest for fundamental theorems of social choice. *Social Choice and Welfare*, 44(3):533–548, 2015.

[26] William S. Zwicker. Introduction to the theory of voting. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 2. Cambridge University Press, 2016.

Justin Kruger
Université Paris-Dauphine
Paris, France
Email: `justin.g.kruger@gmail.com`

Sebastian Schneckenburger
University of Tübingen
Tübingen, Germany
Email: `sebastian.schneckenburger@uni-tuebingen.de`