Strategic Manipulation in Tournament Games*

Allen I.K. Vong†

September 30, 2016

Abstract

I consider the strategic manipulation problem in multistage tournaments. In each stage, players are sorted into groups in which they play pairwise matches against each other. A ranking of players for each group is established according to the match results, and higher ranked players qualify to the next stage. Players prefer qualifying to higher stages. In this setting, a player may potentially profit by exerting zero effort in some matches even when effort exertion is costless. Since such behavior manipulates the tournament, it is desired that full effort exertion is an equilibrium and any equilibrium ranking of qualifying players in each group is immune to manipulation, irrespective of players' strengths. To satisfy these conditions, I show that it is both necessary and sufficient to allow only the top-ranked player from each group to qualify in every stage. Thus, in a tournament with multiple qualifiers in some group, rankings of players can become a noisy indicator of their strengths, while effort cost and heterogeneous prize spread can be of little relevance to players' effort choices.

Keywords: tournament design; contest design; strategic manipulation; subgame perfect implementation; incentive-compatibility.
JEL classification: C7, D7, J7.

*I am indebted to Claudio Mezzetti for his guidance and encouragement throughout the entire project. I particularly thank Johannes Hörner for his support. I also thank the advisory editor, two anonymous referees, Georgy Artemov, David Delacrétaz, Chris Edmond, Simon Loertscher, Ellen Muir, James Nesbit, Jun Xiao, seminar participants at the 27th International Conference on Game Theory in Stony Brook and the 5th World Congress of the Game Theory Society in Maastricht for insightful comments at various stages of the project. This paper supersedes my honors essay at the University of Melbourne.

†Yale University. Email address: iokuan.vong@yale.edu.

1 Introduction

Tournaments are commonly used as an incentive scheme to elicit costly effort from economic agents.[1] However, many tournament designs incentivize shirking behavior from players, *irrespective of any effort cost*. Consider a simple example. Two players, a and b, play against each other in stage 1 of a tournament. The winner plays against c while the loser plays against d in stage 2. The winner of each match in stage 2 qualifies to be a final winner of the tournament. Suppose that when each player plays to win, a always beats b and d, but loses to c. Further, d always beats b and c when he plays to win, while b beats c. Each player prefers qualifying to a higher stage. If they play to win throughout, then a is eliminated in stage 2 and d is the final winner. However, a one-stage deviation by a to shirk and lose against b in stage 1 manipulates the qualification process and makes him the final winner at the expense of d, who is now eliminated in stage 2. Importantly, effort cost is irrelevant in reaching the decision to shirk. A famous real-world example is the 2012 London Olympics, where eight badminton players were disqualified for deliberately losing matches to meet preferred opponents in the next round.

The goal of the present paper is to characterize the set of tournaments immune to such manipulation. This is of interest for three reasons. First, the manipulation problem is in general detrimental to the profitability of organizing a tournament and the reputation of its organizers, and may arise in any economic setting with multistage competition. Second, it is often against a designer's incentive to simply disqualify any player who shirks in any tournament. In particular, it is often the stronger players who profit by shirking, and they shirk to secure their participations in later stages. The matches in later stages that feature stronger players tend to generate larger revenues. Finally, the outcome of a tournament is supposed to reflect the relative strengths of players. If the outcome is manipulated so that a player fails to qualify in an earlier stage than some weaker player, the outcome becomes a noisy indicator of the players' strengths.

To set the stage, I develop a simple model that allows for an *arbitrary* design of a tournament. This contrasts with existing work in the contest design literature, which largely studies the prize spread or assignment of opponents to players given a *specific* design. For example, see Rosen (1986), Moldovanu and Sela (2001), Groh, Moldovanu, Sela, and Sunde (2012) and Fu and Lu (2013). In particular, I build upon the traditional approach in the mathematics literature by defining a tournament as a pair consisting of a set of players and

[1] For example, they are used in sports events, labor promotions and innovation contests. Seminal papers include Lazear and Rosen (1981), Green and Stokey (1983) and Rosen (1986). Konrad (2009) provides a comprehensive survey.

a binary relation on the set which captures the relative strength between any two players.[2] A designer then chooses a design for any given tournament, inducing a *tournament game* – a multistage game with observable effort. Succinctly, in a tournament game, each stage is defined as a partition of the set of active players, while each element in the partition is a *group* of players, within which the players play pairwise matches against all others. The result of each match depends on the efforts exerted by the players and according to the match results, a ranking of players is established for each group. Qualification of each player to a higher stage depends on his rank in a group, subject to a quota. To illustrate, a *round-robin* tournament game has one stage of active play, in which all players belong to the same group. A *single-elimination* tournament game has a sequence of stage partitions, each of which contains a number of groups of two, with the winner in each group qualifying to the next stage. The 2014 FIFA World Cup had eight groups of four teams from which two qualified in stage 1, and played as a single-elimination tournament beginning from stage 2.

More precisely, I characterize the set of *incentive-compatible* tournament games which are tournament games with the following two properties, irrespective of players' strengths. First, to capture sequential rationality, full effort exertion is a *subgame perfect equilibrium*, so that there exists at least one equilibrium with no manipulation. Second, the outcome under *each* subgame perfect equilibrium must coincide with one of the outcomes prescribed by the *full-effort social choice correspondence*, which maps the set of players, their relative strengths and a given design to the set of possible outcomes under which the strongest players from a group qualify and amongst these qualifying players a stronger player ranks higher. Put simply, this protects the outcome from manipulation even when a deviation away from full effort exertion to another equilibrium takes place. Plainly, the main result we are going to prove is the follwing: *Every incentive-compatible tournament game allows only the top-ranked player from each group to qualify.* In the proof, we shall see that in tournament games with multiple qualifiers in some group, some player can often profit by shirking to qualify with a lower rank. More precisely, in *every* tournament game with multiple qualifiers in some group, there always exists an equilibrium with an outcome different from those prescribed by the full-effort social choice correspondence.

This paper differs from related work by providing an analytical *game-theoretic foundation* and allowing for an *arbitrary design*. For instance, Dagaev and Sonin (2013) borrows tools from social choice theory to show that multiple qualifier systems, where players compete in several local tournaments to qualify to international tournaments, are manipulable. On the other hand, Pauly (2014) uses a computer-assisted proof to show that the designs of the 2012 Olympics Badminton and the 2014 FIFA World Cup give rise to manipulation.

[2]For example, see Rubinstein (1980). Laslier (1997) provides an extensive survey.

2 Model

Let N be a finite set of players. A tournament is a pair $(N, \to)$, where $\to$ is a complete, asymmetric and irreflexive binary relation on N, so that $i \to j$ if i is *(relatively) stronger than* j. Given N, a designer fixes a *design* D_N which induces a tournament as a multistage game, referred to as a *tournament game*. A design consists of four components:

1. the number of stages and groups in each stage, as well as the size, a distinct label and a qualification quota of each group;
2. a partition P^1 of the set N, where each element $G \in P^1$ is a group in stage 1;
3. a sorting rule f^t for each stage $t > 1$ that sorts players qualifying from stage $t-1$ into groups in stage t;
4. a tie-breaker $\succ_{\text{tie}}$, which is a strict linear order on N.

The order $\to$ on N is known to the players and the designer only *after* the design is fixed. A player is *active* before he is eliminated. In every stage t, each active player is sorted into some group G. Within G, every player i plays a *match* against each opponent j, where each *costlessly* chooses an effort level from $[0, 1]$ against his opponent. The player who exerts a higher effort wins the match. When both exert the same effort, then i wins if $i \to j$. Let $w_i(G) \equiv |\{\iota \in G : i \text{ wins against } \iota\}|$ denote i's number of wins in G. Define a ranking of players $\succ_G$ on G such that $i \succ_G j$, or *i ranks higher than j in G*, if $w_i(G) > w_j(G)$, or if $w_i(G) = w_j(G)$ and $i \succ_{\text{tie}} j$. Also let $w_i^{FE}(G) \equiv |\{\iota \in G : i \to \iota\}|$ denote i's number of wins in G if all players exert the same effort (e.g., full effort, or effort 1).

The q_G highest ranked players in G by $\succ_G$ qualify from G to the next stage. Let $|G| \geq 2$ and $1 \leq q_G \leq |G| - 1$, so that at least one player is eliminated and at least one qualifies from each group. In the next stage, active players are sorted into groups by a sorting rule. A sorting rule f^t partitions the set of active players in stage t according to the label of the groups they qualify from in stage $t-1$ and their ranks in the groups, with the partition denoted by P^t such that each $G \in P^t$ is a group. The game ends when players from stage $T \geq 1$ qualify to stage $T+1$, where no sorting of players and no match take place.

Efforts are observed at the end of each stage and there is perfect recall. Let H^t be the set of histories in the beginning of stage t, which contains a typical element $h^1 = \emptyset$ or $h^t = (a^1, \ldots, a^{t-1})$ for $t \geq 2$, where $\emptyset$ denotes the null history and a^t is the vector of efforts exerted by players in stage t. Let $A_i(h^t)$ be player i's set of feasible effort choices in stage t when the history is h^t. A pure strategy by i is a sequence of maps $(s_i^t)_{t=1}^T$, where each $s_i^t : H^t \to \cup_{h^t \in H^t} A_i(h^t)$, so that $s_i^t(h^t) \in A_i(h^t)$ for every h^t. Each i chooses a strategy

$s_i = (s_i^t)_{t=1}^T$ to maximize his payoff $u_i : H^{T+1} \to \{1, \ldots, T+1\}$ defined by $u_i(s_i, s_{-i}) = t$, given the strategy s_{-i} by all players except i, where t is the highest stage player i qualifies to. Of particular interest is the profile s^*, where players exert full effort throughout (i.e., after all histories).

An *outcome* of a tournament game is a collection $\mathcal{O}_{((N,\to),D_N)} = ((G, (\succ_G))_{G \in P^t})_{t=1}^T$. Given $(N, \to)$, define an *outcome function* g as $s \mapsto g(s) = \mathcal{O}_{((N,\to),D_N)}$ for any $s \in H^{T+1}$. A *social choice correspondence* (henceforth, SCC) F maps $((N, \to), D_N)$ into a set of outcomes $\mathcal{O}_{((N,\to),D_N)} \in F((N, \to), D_N)$. Let F^{FE} be the *full-effort* SCC where $F^{FE}((N, \to), D_N)$ is the set of outcomes such that in each stage $t = 1, \ldots, T$, for two players i, j who qualify and any player k who is eliminated from some group $G \in P^t$, $i \succ_G j \succ_G k$ if and only if $w_i^{FE}(G) \geq w_j^{FE}(G) \geq w_k^{FE}(G)$ and the tie-breaker $i \succ_{\text{tie}} j \succ_{\text{tie}} k$ applies whenever equality holds.

A pair (F, s) of a SCC F and a strategy $s \in H^{T+1}$ is *subgame perfect implementable* if s is a subgame perfect equilibrium (henceforth, SPE) and $g(s') \in F((N, \to), D_N)$ for *any* SPE $s' \in H^{T+1}$. There is at least one SPE because a tournament game is a finite game of perfect information.

A tournament game is *incentive-compatible* (henceforth, IC) if (F^{FE}, s^*) is subgame perfect implementable for every order $\to$ on N. Intuitively, the universal quantification is desirable because the relative strengths are unknown before a design is fixed. Examples 2 and 3 below show that the two requirements in the implementation of (F^{FE}, s^*) are independent.

Remark 1. The fact that effort is costless and that players derive a unit marginal payoff upon qualifying over each stage, albeit stylized, serves to emphasize that effort cost and prize spread can be of little relevance to players' effort choices when a tournament is not IC. Moreover, pairwise matching stands in contrast to a more general assumption where multiple players can possibly be matched simultaneously. For our purpose, pairwise matching delivers an advantage by its clear application: a player often chooses a specific opponent against whom he shirks, and pairwise matching allows us to capture such action.

3 Incentive-Compatible Tournament Games

In this section, we formally state the main theorem. Perhaps surprisingly, as far as IC is concerned, all components in a design except the qualification quota in each group are irrelevant.

Theorem 1. *Fix a tournament $(N, \to)$. A tournament game induced by a design D_N is IC if and only if the design designates that $q_G = 1$ for every group G in any stage $t \leq T$.*

Before proceeding to the proof, we study a few examples to illustrate the model and to build intuition underlying the result.

Example 1 illustrates a design which, with multiple qualifiers allowed in some group, induces a tournament game where s^* fails to be a SPE. Denote by $G[r]$ the r^{th} ranked player in a group G.

Example 1. Consider a tournament $(N, \to)$, where $N = \{a, b, c, d, e, f\}$ and

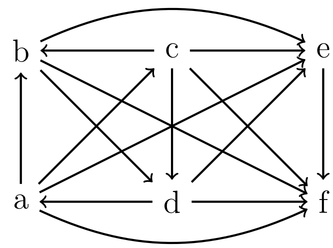

.

The design D_N is as follows. The first stage partition is $P^1 = \{A, B\}$, where $A = \{a, b, e\}$ and $B = \{c, d, f\}$, with $q_A = q_B = 2$. Then $P^2 = \{C, D\}$, where $C = \{A[1], B[2]\}$ and $D = \{A[2], B[1]\}$, with $q_C = q_D = 1$. The players $C[1]$ and $D[1]$ are the final winners. Consider the outcome under s^*. In stage 1, $a \succ_A b \succ_A e$ and $c \succ_B d \succ_B f$. In stage 2, $C = \{a, d\}$ and $D = \{b, c\}$, with $d \succ_C a$ and $c \succ_D b$. The final winners are c and d. If a unilaterally deviates to shirk against b in stage 1, then the ranking in A becomes $b \succ_A a \succ_A e$, so that $C = \{b, d\}$ and $D = \{a, c\}$ in stage 2. Because $a \succ_D c$, a qualifies as a final winner and profits from the deviation.

The key insight from the example is that we can always find a relation $\to$ on N such that s^* fails to be a SPE whenever a player can "choose" which group to qualify to. Nonetheless, this should not be confused with the statement that one can always find a relation $\to$ on N such that s^* fails to be a SPE whenever multiple qualifiers are allowed in some group. In particular, Example 2 shows that the latter statement is false. Moreover, the example also shows that while s^* is a SPE, there exists another SPE s inducing an outcome $g(s)$ that does not agree with any outcome prescribed by the full-effort SCC.

Example 2. Consider a tournament $(N, \to)$ where $N = \{a, b, c\}$, $a \to b$, $b \to c$ and $a \to c$. The design D_N is as follows. The first stage partition is $P^1 = \{A\}$, where $A = \{a, b, c\}$ and $q_A = 2$. Then $P^2 = \{B\}$, where $B = \{A[1], A[2]\}$ and $q_B = 1$. $B[1]$ is the final winner. The outcome implemented by the SCC F^{FE} is such that $a \succ_A b \succ_A c$; $B = \{a, b\}$ and $a \succ_B b$, implying that a is the final winner. It is easy to verify that s^* is a SPE. Further, consider a strategy profile s equivalent to s^* except that a shirks against b in stage 1. Under s, $b \succ_A a \succ_A c$. Then $B = \{a, b\}$, $a \succ_B b$ and a is the final winner. Being the final winner, a

has no profitable deviation. Neither does b, because by shirking against a in either stage he cannot change his payoff, and by shirking against c in stage 1, he can possibly be eliminated a stage earlier, depending on the specified tie-breaker. Finally, it is clear that c has no profitable deviation. Thus there exists a SPE s such that $g(s) \notin F^{FE}((N, \rightarrow), D_N)$.

Conversely, Example 3 below illustrates a setting where $g(s) \in F^{FE}((N, \rightarrow), D_N)$ for any SPE s, but s^* is not a SPE in the tournament game. Together with Example 2, this shows that the requirements of s^* being a SPE and $g(s) \in F^{FE}((N, \rightarrow), D_N)$ for every SPE s are independent.

Example 3. Consider a tournament $(N, \rightarrow)$, where $N = \{a, b, c, d, e, f, g\}$ and

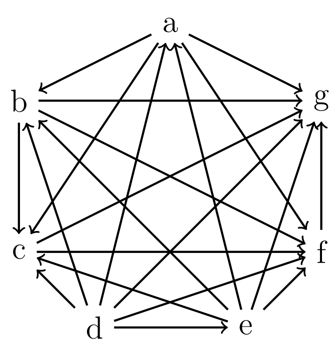

.

Let the design D_N designate that $P^1 = \{A, B, C\}$ where $A = \{a, b, c\}$ with $q_A = 2$, $B = \{d, e\}$ with $q_B = 1$ and $C = \{f, g\}$ with $q_C = 1$. Then $P^2 = \{D, E\}$, where $D = \{A[1], B[1]\}$ and $E = \{A[2], C[1]\}$. The players $D[1]$ and $E[1]$ are the final winners. Finally, suppose that the tie-breaker follows $c \succ_{\text{tie}} b \succ_{\text{tie}} a \succ_{\text{tie}} d \succ_{\text{tie}} e \succ_{\text{tie}} f \succ_{\text{tie}} g$. By construction of $\rightarrow$, by exerting full effort in group E, $A[2]$ ultimately becomes a final winner, while $A[1]$ is eliminated from group D in stage 2 as long as $B[1]$ exerts full effort. Moreover, in *any* SPE, both a and b never rank the lowest in A and are eliminated, for a always qualifies by exerting full effort in every match in A, while b always qualifies by exerting full effort against c. So in equilibrium, either $a \succ_A b \succ_A c$ or $b \succ_A a \succ_A c$. If the latter is true, then b must have 2 wins, a has 1, and c has 0, by construction of $\succ_{\text{tie}}$. Thus a must have shirked against b. But then b can be strictly better off by shirking against a and qualify as $A[2]$, contradicting subgame perfection. So it must be true that $a \succ_A b \succ_A c$, and a has 2 wins, b has 1 and c has 0 by construction of $\succ_{\text{tie}}$. In particular, b must have shirked against a, for otherwise a could shirk against him and profit by qualifying as $A[2]$. On the other hand, it should be clear that the stronger player in each group B, C, D and E would exert full effort in equilibrium. Thus $g(s) \in F^{FE}((N, \rightarrow), D_N)$ in any SPE s. But s^* is clearly not a SPE in the induced tournament game, for a would profit by unilaterally deviating to shirk against b in A to qualify as $A[2]$.

Next, recall that F^{FE} imposes a restriction on the rankings of *qualifying* players instead of *all* players in each group. In particular, Example 4 shows that the latter requirement is so restrictive that even a simple round-robin tournament is not subgame perfect implementable.

Example 4. Consider a tournament $(N, \rightarrow)$, where $N = \{a, b, c, d, e, f\}$ and

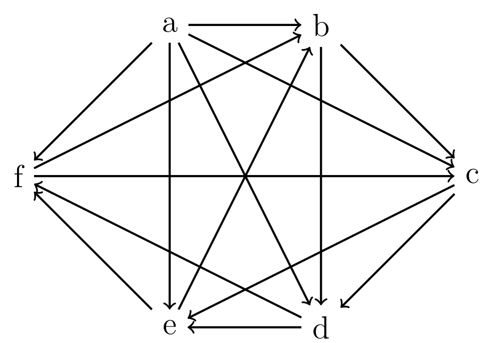

.

The design is one of *round-robin*: $P^1 = \{G\}$ where $G = \{a, b, c, d, e, f\}$, and $G[1]$ is the final winner. Consider a strategy profile s equivalent to s^* except that a shirks against f in G. The tie-breaker satisfies $a \succ_{\text{tie}} b \succ_{\text{tie}} c \succ_{\text{tie}} d \succ_{\text{tie}} e \succ_{\text{tie}} f$. It is straightforward to check that s is a SPE. In particular, under s we have $(w_a, w_b, w_c, w_d, w_e, w_f) = (4, 2, 2, 2, 2, 3)$, where $w_i \equiv w_i(G)$ for each player i. As a result, f ranks strictly above all other players except a in G. Consider a social choice function F_*^{FE} that requires that the ranking of all players, as opposed to only the winner as does F^{FE}, are equal to the ranking under full effort. Then F_*^{FE} designates that $i \succ_G f$ for every $i \in N \backslash \{f\}$, so $g(s) \notin F^{FE}((N, \rightarrow), D_N)$. Intuitively, because a is far stronger than his opponents in the group, he is indifferent between winning all matches or winning all except losing to f. Losing to f, however, manipulates the ranking.

Finally, one key assumption underlying Theorem 1 is that the sorting rule in each stage relies on the ranks of each qualified player instead of their numbers of wins in the previous stage. Indeed, the latter measure provides an immediate medium for a player to manipulate his standing and manipulating behavior may arise even when the qualification quota is one in each group as shown Example 5 below.

Example 5. Fix some stage in an arbitrary tournament game. Suppose that if all players exert full effort then player i ranks first with w wins in some group G the current stage, qualifies to group A in the next stage and is eliminated. Suppose also that no other player in G has more than $w - 3$ wins in the current stage, so i would still qualify if he lost one match in G. Suppose that when i ranks first in G with $w - 1$ wins he qualifies to group B in the next stage and then ranks first in B as long as he exerts full effort. Then exerting full effort after all histories is not a SPE.

We are now ready to prove the theorem.

Proof of Theorem 1. Fix a tournament $(N, \to)$ and a design D_N where $q_G = 1$ for every group $G \in P^t$ in every stage $t \leq T$. Since every stage begins with a proper subgame, it suffices to show that the restriction $s^*|_{h^t}$ of s^* from the beginning of each stage t for each history h^t yields a *Nash equilibrium* in the induced tournament game. This trivially holds for any final winner under s^*. This is also true for a player j who is eliminated in some group G in any stage t under $s^*|_{h^t}$ given any history $h^t \in H^t$, because any deviation $s_j|_{h^t}$ would see him remain eliminated in stage t. By the same token, the player i who qualifies from G and is eliminated in a stage $t' > t$ under s^* has no profitable deviation from $s_i^*|_{h^{t'}}$ given a history $h^{t'}$. Moreover, any deviation by i from $s_i^*|_{h^t}$ in any stage $\tau = t, \ldots, t' - 1$ is not profitable, because by shirking, he may no longer rank first and is eliminated in stage τ, or he remains qualifying with a first rank to the same group in stage $\tau + 1$ and thus plays against the same opponents in subsequent stages until his elimination in stage t'. Thus s^* is a SPE. To show that $g(s) \in F^{FE}((N, \to), D_N)$ for any SPE s, observe that the set of outcomes designated by F^{FE} are those where i qualifies from G if and only if $i \in \text{argmax}_{j \in G}\, w_j^{FE}(G)$ and $i \succ_{\text{tie}} k$ for any other maximizer k. If $g(s) \notin F^{FE}((N, \to), D_N)$ for some SPE s, then there must exist a group G in some stage t after some history h^t, and two players $i, j \in G$ so that $i \succ_G j$ but $j \in \text{argmax}_{l \in G}\, w_l^{FE}(G)$ and $j \succ_{\text{tie}} k$ for any other maximizer k. Because $q_G = 1$, j is eliminated in stage t. But the fact that j is a maximizer implies that if j deviates to play $s_j^*|_{h^t}$, then $j \succ_G \iota$ for any other player ι in G given any $s_{-j}|_{h^t}$ and therefore j qualifies, contradicting that s is a SPE.

Conversely, it suffices to show for each design D_N with $q_G > 1$ in a group G in some stage t that, for some order $\to$ on N, either s^* is not a SPE or that $g(s) \notin F^{FE}((N, \to), D_N)$ for some SPE s. Fix a history h^t, after which $G = \{1, 2, \ldots, g\}$. Let the order $\to$ on N designate that $1 \to j$ for each $j \in G \backslash \{1\}$ and $2 \to j$ for each $j \in G \backslash \{1, 2\}$, and also that for any two players $i, j \in G$, $i \to j$ if $i < j$. Thus $i \succ_G j$ if $i < j$ under $s^*|_{h^t}$, and both 1 and 2 qualify from G. Upon qualification, two possibilities arise, depending on the sorting rule f^{t+1}. The first possibility entails that $1, 2 \in A$ for some group A in stage $t + 1$. Let s^* be a SPE, for otherwise the claim is proved in this case. Consider a one-stage deviation s from s^* in stage t where 1 shirks against 2 in G. The rankings of all groups in P^t are unchanged, except that the ranks of 1 and 2 switch in G so that $g(s) \notin F^{FE}((N, \to), D_N)$. IC fails because s is a SPE since first, the stage partitions $(P^t)_{t=1}^T$ are the same under s and s^*, which implies that players face the same opponents in each stage under both profiles and second, s^* is a SPE. It remains to consider the second possibility, where 1 and 2 are sorted into two separate groups, say B and C respectively, in stage $t + 1$. Further let the order $\to$ on N designate that, $i \to 1$ for each $i \in B \backslash \{1\}$ and $1 \to j$ for each $j \in C$. Clearly, 1 ranks the lowest in B

and is eliminated in stage $t+1$. However, s^* is not a SPE because a one-stage deviation by 1 to shirk against 2 in G would allow 1 to qualify to C in stage $t+1$, in which he ranks the first and qualifies further. □

4 Final Comment

The result has the flavor of an impossibility result by implying that many tournament games in practice are not IC. Nonetheless, in practice, regulations that allow only one qualifier from each group may not be desirable. The appeal of IC may therefore be put into question. Imagine an unknown tennis player being drawn against Roger Federer in the same group in the first stage of a single-elimination tournament. The tournament game is IC, but is perhaps too harsh on the athletes and their fans. Of course, there are many factors other than IC, for instance time constraint and the number of matches, that a designer needs to take into account when designing a tournament. To be clear, the aim of the present paper is *not* to provide a general selection criterion of tournament designs. Instead it seeks to provide a baseline framework for future work. For instance, the current solution concept is *binary*: a tournament game is either IC or not. Devising a method to quantify IC would allow for the study of trade-offs between IC and other potentially desirable factors.

References

D. Dagaev and C. Sonin. Winning by losing: Incentive incompatibility in multiple qualifiers. *Discussion paper series*, Centre for Economic Policy Research, 2013.

Q. Fu and J. Lu. The optimal multi-stage contest. *Economic Theory*, 51(2):351–82, 2013.

J. Green and N. Stokey. A comparison of tournaments and contracts. *Journal of Political Economy*, 91(3):349–64, 1983.

C. Groh, B. Moldovanu, A. Sela, and U. Sunde. Optimal seedings in elimination tournaments. *Economic Theory*, 49(1):59–80, 2012.

K. Konrad. *Strategy and Dynamics in Contests*. Oxford University Press, New York, 2009.

J.-F. Laslier. *Tournament Solutions and Majority Voting*, volume 7 of *Studies in Economic Theory*. Springer, 1st edition, 1997.

E. Lazear and S. Rosen. Rank-order tournaments as optimum labor contracts. *Journal of Political Economy*, 89(5):841–64, 1981.

B. Moldovanu and A. Sela. The optimal allocation of prizes in contests. *American Economic Review*, 91:542–58, 2001.

M. Pauly. Can strategizing in round-robin subtournaments be avoided? *Social Choice and Welfare*, 43(1):29–46, 2014.

S. Rosen. Prizes and incentives in elimination tournaments. *American Economic Review*, 76 (4):701–15, 1986.

A. Rubinstein. Ranking the participants in a tournament. *Journal of the Society of Industrial and Applied Mathematics*, 38:108–11, 1980.