The Scandinavian
Journal of Economics

Scand. J. of Economics 124(4), 1087–1116, 2022
DOI: 10.1111/sjoe.12476

Peer interactions and performance in a high-skilled labour market*

Matteo Pazzona

Brunel University London, Middlesex UB8 3PH, UK
matteo.pazzona@brunel.ac.uk

Abstract

It is not clear whether interactions among superstar employees lead to an increase in productivity. Such interactions are relatively rare, and measuring productivity is challenging. In this paper, it is suggested that these difficulties can be overcome by analysing changes in the performance of elite National Basketball Association (NBA) players who participate in the Olympic Games. By using advanced individual performance measures, the study finds that these athletes experience an increase in performance of 7.1 percent in the season after the Games, compared with similar non-Olympic athletes. The sharp discontinuity in peer quality experienced by the players is the most likely explanation for this increase.

Keywords: High-skilled labour market; individual performance; National Basketball Association; peer interactions and quality

JEL classification: *J*01; *J*24

1. Introduction

This paper aims to answer the following question: do even the most highly skilled workers benefit from working with similarly talented colleagues? Such a question is important because these workers expand the frontiers of knowledge in their professions. In turn, superstar workers can affect the performance of lower-skilled co-workers. Understanding the determinants of highly skilled workers' performance is economically relevant. However, such a theory is not easily tested: real-life examples of the interactions between highly skilled workers are less common than those for other skill-type combinations (high with low or low with low). Moreover, it is challenging to measure the contribution of a worker to the success of their firm. I try to overcome these difficulties in the context of a

*I would like to thank two anonymous referees. I received useful comments from Matias Busso, Alejandro Corvalan, Jessica Gagete-Miranda, and Romain Gauriot. I would also like to acknowledge participants at the COMPIE conference, the European Economic Association annual conference, and the SESM conference. I finally thank participants in the seminars at the Diego Portales University, the University of Kent, the University of Reading, and the University of Santiago.

very highly skilled labour market: professional basketball players. I do so using the positive shock in peer quality experienced by National Basketball Association (NBA) players participating in the Olympic Games.

Since 1992, US Olympic players have been considered the elite of professional basketball and include sporting legends such as Michael Jordan and Kobe Bryant. During the summer of each Games, selected players spend several weeks training and playing together.[1] Does the Olympic experience lead to better performance when they return to their NBA teams? To answer this question, I compare the change in productivity between the players who went to the Olympics and those who did not go, using a difference-in-difference strategy. To measure the impact of a player on the success of the team, I take advantage of recent developments in individual advanced performance measures, which make it possible to precisely identify a player's contribution. Given that selection for treatment is not random, I employ the propensity score of selection for the Olympic team to calculate the kernel weights, which are then applied to the difference-in-difference analysis. I show that, in the season after each Games, Olympic players increase their player efficiency rating (PER) – the preferred measure of performance – by 7.1 percent. The baseline findings are confirmed when several potential threats to the identification strategy are taken into account. For example, I consider three placebo treatments, as well as different ways in which the propensity score is calculated. The robustness exercises all support the hypothesis that the increase in performance must be considered as causally linked to participation in the Games. Which channels are most likely to explain this result? I argue that by going to the Games, Olympic players experience a positive shock in peer quality, while the control athletes do not. During the regular NBA season (October/November to April), these superstars compete alongside players who are, on average, less talented than them. Even though some teams are better than others, the salary cap rules avoid an excessive concentration of talent within any one team in the NBA. When these players join the Olympic team, however, they compete alongside elite athletes, who have an average PER that is about 67 percent higher than that of their NBA teams. This difference in the ability of teammates – a significant positive peer shock – can help to explain the increase in performance. Based on this premise, I regressed the change in performance of a player between the periods before and after the Games on to the difference in teammate quality between the Olympic team and the player's original NBA team. I found evidence that players with a greater peer discontinuity registered the greatest increase in performance. Nevertheless, there might be other explanations for the increase in productivity, as well

[1] Players usually spend between one and one and a half months together.

as peer effects. For example, I also considered the role of discontinuity in the quality of coaches and opponents, but was unable to identify a more convincing explanation than peer effects.

The return of the Olympic athletes to their NBA teams allowed me to determine whether their presence led to trickle-down effects on lower-skilled teammates the season after the Games. To test this, I ran a difference-in-difference regression, considering other players from the same NBA team as the Olympic players. This exercise revealed the absence of positive externalities on lower-skilled teammates. I also investigated whether players with skill levels above and below the median were affected heterogeneously, but I found no effect.

By providing clear evidence of the benefits of interaction between highly skilled workers, I make a relevant contribution to the existing literature on peer effects in the workplace.[2] It is typical for this literature to make a distinction between the learning effect and the motivation effect (Guryan et al., 2009; Cornelissen et al., 2017). The former considers how a worker learns the best way of performing a task from their co-workers. Studies that aim to isolate the learning effect have mainly focused on highly skilled jobs, as they are typically non-repetitive and require a substantial degree of creativity and sophistication. This literature has mainly studied teachers and scientists, finding mixed evidence (Jackson and Bruegmann, 2009; Azoulay et al., 2010; Waldinger, 2011). The motivation effect refers to the fact that a worker is motivated when their co-workers are doing well. The literature on the motivation effect is based mainly on lower-skilled workers, who are typically employed in occupations that are repetitive and allow direct observation of outputs (Cornelissen et al., 2017). The literature displays some consensus about a positive effect (Falk and Ichino, 2006; Mas and Moretti, 2009; Bandiera et al., 2010; Kaur et al., 2010). Nevertheless, it is often difficult to make a clear distinction between the learning and motivation effects in the workplace, especially when both are acting simultaneously (Gould and Winter, 2009; Guryan et al., 2009; Hickman and Metz, 2018). In the case analysed in this paper, superstars learn from their Olympic teammates who have specialized in other tasks and thus benefit from knowledge spillover. At the same time, motivation effects can also be triggered by increased confidence. Playing all summer with high-calibre athletes can boost a player's confidence that he can compete at the highest levels. In Section 5, I provide some exercises to distinguish between these factors.

[2]Given that the focus of this study is on peer interaction in the workplace, the education literature is not exhaustively mentioned. This branch of the literature is richer than that on peer effects in the workplace. For a review of peer effects in education, refer to Sacerdote (2011).

My work is also related to, and complements, existing studies of peer effects in sports (Guryan et al., 2009; Depken and Haglund, 2011; Yamane and Hayashi, 2015; Emerson and Hill, 2018; Jiang, 2020). However, when compared with this literature, my study has some unique features. First of all, in my setting, athletes experience a peer shock that lasts for more than a month, when they continually train and interact with the best in their field. In the other studies, the period of peer exposure is much more limited. Additionally, my analysis focuses on the impact on performance in the seasons after the peer interactions took place. This allows the medium- and long-run effects to be evaluated, while the existing literature focuses mainly on the immediate impact on performance.

This paper also contributes to the literature on the effect of highly performing workers on lower-skilled co-workers (Brown, 2011; Agrawal et al., 2017; Serafinelli, 2019), which has found mixed evidence. In the setting studied here, the athletes returned to their original team after the interactions with their peers, allowing me to study the role of trickle-down effects.[3] Unlike the existing literature, this work employs individual performance statistics, enabling me to better evaluate the impact of the workers (players) on the success of their firms (teams). To the best of my knowledge, this is the first time that such performance measures have been used to evaluate productivity in this context.

Alongside the existing academic literature, there is much anecdotal evidence that supports the peer-effect interpretation. In a recent interview, referring to his experience in the 2008 Olympics with Kobe Bryant, Dwayne Wade said: "[W]ith the Olympics, you see a guy daily [. . .]. You get to see his work ethics, you get to see, you get to be around him to hear his knowledge of the game, you get to play with a guy. You are in the trenches with a guy."[4,5]

As mentioned, peer effects might not be the only determinant of the increase in each player's performance. The results of this study might also be consistent with the existence of team incentives Hamilton et al., 2003; Babcock et al., 2015, training (Becker, 2009; De Grip and Sauermann, 2012), the identity of the manager running the team (Lazear et al., 2015),

[3]Ichniowski and Preston (2014) have shown that the presence of players from elite clubs positively affected the performance of national teams in soccer. In a sense, my study explores the inverse setting, where peer effects move to the original teams from the national ones.

[4]See the article, "Kobe Chronicles: Dwyane Wade and Bryant built mutual respect on Team USA", by L. Thiry, in the *Los Angeles Times*, 11 April 2016, https://www.latimes.com/sports/lakers/la-sp-lakers-kobe-chronicles-dwyane-wade-kobe-bryant-20160410-story.html.

[5]Michael Jordan has repeatedly said that the greatest game he has ever played was a scrimmage in the summer of 1992 between teammates in preparation for the Olympics (see the ESPN article by M. Adams, The Dream Team scrimmage in Monte Carlo, https://www.espn.com/blog/statsinfo/post/_/id/133080/the-scrimmage-in-monte-carlo).

or the quality of the team's opponents. I explore such possibilities in the section on channels.[6]

What happens to NBA players has important implications for other labour markets where highly skilled workers are distributed across different firms (Kahn, 2000; Rosen and Sanderson, 2001). For example, an analogous approach could be used to evaluate changes in the academic productivity of scholars who participate in exchange programmes. More generally, a similar setting could be found whenever there is a significant discontinuity in the quality of peers. From the perspective of firms, it might be a good investment to promote collaboration and create a kind of "all-star" entity/firm in which only the most talented employees take part. However, firms should be aware that lower-skilled co-workers would probably not benefit.

The paper is organized as follows. Section 2 describes the selection process of the Olympic Games and the identification of the treated and control groups. It also presents the results of the propensity score exercises employed to calculate the kernel weights. Section 3 presents the econometric model and gives the baseline results alongside the table showing the exercises, which take into account possible challenges to the identification. Section 4 presents further tests alongside the results of the dynamics and heterogeneous treatment. Section 5 then explores the channels that might be responsible for the increase in performance. Section 6 investigates the presence of trickle-down effects on the Olympic players' teammates. Finally, I conclude in Section 7 and discuss my findings.

2. Context, background, and data

2.1. Context and performance measures

The US national basketball team is by far the most successful in the history of the sport at the summer Olympic Games. Out of the 19 Games that the United States has taken part in, it has won the gold medal 16 times, the silver medal once and the bronze twice. This is despite professional NBA players only being allowed to play since the 1992 Games in Barcelona. From that point, the Olympic athletes have always been the elite of the NBA, and thus of the world. For example, the 1992 team featured Hall of Fame players of the calibre of Michael Jordan, Magic Johnson, and Larry Bird. How were these players selected? On what basis did the selection

[6]A relatively recent paper on professional ice hockey (Cairney et al., 2015) documented a decrease in the performance of professional athletes after the winter Olympic Games. However, the Winter Olympic Games are held in the middle of the regular season whereas the Summer Games take place between two different seasons.

committee choose them? Although there are no official documents that explicitly set out the selection criteria, it is most likely that the selection was made by considering a range of their characteristics, including the player's ability, indicated by his performance in the seasons before the Olympic Games. Other characteristics, such as a player's career trajectory, age, experience, and position on the court might also have been important in influencing the selection committee. I go into more detail about these variables in Section 2.2, and here I focus first on the key variable: player performance.

Defining attainment is not easy in sport, especially in team sports such as basketball. The marginal contribution of a player in a team setting is the result of complex dynamics, including productivity spillovers between teammates (Kuehn, 2017). A player can have a strong impact on his team, even though his individual statistics fail to capture that contribution (Oliver, 2004). This problem is also found in other types of team-related jobs, where wages are only an imperfect measure. In basketball, as previously in baseball, advanced individual performance statistics have been developed to capture the whole contribution of a player to the success of his team.

One of the most frequently used advanced performance measures – and the preferred one in this paper – is the PER, which synthesizes a player's different accomplishments in a single measure.[7] The PER belongs to the family of linear weights (Hollinger and Hollinger, 2005), in which different statistics are added or subtracted, according to particular weights decided by the developer of the metric (Kubatko et al., 2007). The positive accomplishments of basketball players include points, assists, and rebounds. These are added. Negative accomplishments are subtracted, and include turnovers and personal fouls. The PER is a minute-by-minute measure of a player's performance, which makes it possible to compare athletes with different playing times. It can be further adjusted by the team's pace – in other words, its average possession in that season. This means that the measure does not penalize players in teams that have a slower rhythm. The PER's league average is set at the same level each year, which makes it possible to compare the performance of individuals in different seasons.[8] The PER has been chosen over other measures because it is intuitive and can also be applied to non-sporting contexts. Moreover, as we see in the next section, in the competition between different performance measures, the PER is found to be the best predictor of selection for the Olympic team. Alternative measures based on Plus/Minus statistics and Win Shares

[7] Performance measures, along with other data, have been retrieved from the website https://www.basketball-reference.com.

[8] Michael Jordan and Lebron James are the players with the highest averages throughout their careers. Not surprisingly, they are considered to be among the best players in NBA history.

will be employed to improve robustness. Section A of the Online Appendix provides an exhaustive description of these measures, although the results are consistent using all the performance measures.

2.2. Treated individuals and controls

The treated individuals are those players who participated in the summer Olympic Games from 1992 to 2016. The few players selected from the National Collegiate Athletic Association (NCAA) and those who did not play the season immediately after the Games were excluded.[9] The list of players can be seen in Table 1. In total, the analysis is based on 79 treated players, who represent the elite of professional basketball. The average PER of these athletes in the season before the Olympic Games was 21.75. In contrast, the average for all the other NBA players with a US passport for the same period was 13.3. This means that Olympic players are 1.53 standard deviations better than the average US NBA player.

Selection for the Olympic teams is not a random process – only superstars are selected. The main identification challenge for me – required to evaluate changes in performance – was to find a suitable control group that would have the same trajectory as the selected players in the absence of treatment. To do so, I mimicked the selection process by matching treated and control units according to their propensity score (i.e., their conditional probability of participating in the Olympic Games). More formally, I estimated the propensity score, $p(X_i)$, using the probability model $Pr(S_i = 1|X_i) = F\{h(X_i)\}$, where $S_i = \{0, 1\}$ depends on whether the player i was selected – and participated. X_i represents the pre-treatment characteristics that are likely to affect participation. Because I used probit to estimate the probability, $F(\cdot)$ is the normal distribution and $h(X_i)$ is the function of observable variables. I consider the two seasons before the summer of the Olympic Games as the pre-treatment period.[10] For example, for the 1992 Games, I consider the seasons 1990–91 and 1991–92. The main variable that affects the selection is the quality of the player, which is proxied by his PER. I consider the average PER in the two seasons before the treatment. However, the decision to choose a player might have been affected not only by his performance level, but also by his career trajectory. To capture this feature, I considered the percentage change in performance between the penultimate and last seasons before the

[9]For example, in 2012 Anthony Davis was selected directly from the NCAA. Magic Johnson did not play in the NBA after the 1992 Games.
[10]The choice of two years was made to take into consideration that each extra year means excluding from the sample players who entered the league shortly before the Olympics, thus reducing the number of potential treated and controls.

Table 1. List of Olympic athletes

1992	1996	2000	2004	2008	2012	2016
Barkley, C.	Barkley, C.	Abdur-Rahim, S.	Anthony, C.	Anthony, C.	Anthony, C.	Anthony, C.
Bird, L.**	Hardaway, A.	Allen, R.	Boozer, C.	Boozer, C.	Bryant, K.	Barnes, H.
Drexler, C.	Hill, G.	Baker, V.	Duncan, T.	Bosh, C.	Chandler, T.	Butler, J.
Ewing, P.	Malone, K.	Carter, V.	Iverson, A.	Bryant, K.	Davis, A.*	Cousins, D.
Johnson, M.***	Miller, R.	Garnett, K.	James, L.	Howard, D.	Durant, K.	Derozan, D.
Jordan, M.	Olajuwon, H.	Hardaway, T.	Jefferson, R.	James, L.	Harden, J.	Durant, K.
Laettner, C.*	O'Neal, S.	Houston, A.	Marbury, S.	Kidd, J.	Iguodala, A.	George, P.
Malone, K.	Payton, G.	Kidd, J.	Marion, S.	Paul, C.	James, L.	Green, D.
Mullin, C.	Pippen, S.	McDyess, A.	Odom, L.	Prince, T.	Love, K.	Irving, K.
Pippen, S.	Richmond, M.	Mourning, A.	Okafor, E.*	Redd, M.	Paul, C.	Jordan, D.
Robinson, D.	Robinson, D.	Payton, G.	Stoudemire, A.	Wade, D.	Westbrook, R.	Lowry, K.
Stockton, J.	Stockton, J.	Smith, S.	Wade, D.	Williams, D.	Williams, D.	Thompson, K.

Notes: The table reports the list of players who participated in the Olympic teams in the seven editions under consideration. Players with a star sign are excluded from the analysis: * refers to the athletes who did not play in the NBA in the season before the Games; ** refers to the athletes who did not play in the NBA after the Games; *** refers to the athletes who did not play in the seasons before and after the Games.

Olympics. By controlling for the trend, I was able to improve the quality of the controls, which acted as counterfactuals for the Olympic players in the absence of treatment. Continuing with the career trajectory argument, selection is likely to be influenced by the number of seasons played by an athlete; experienced players are likely to manage pressure better than less experienced ones. However, age is another crucial factor that the selection committee might consider: a young player is less likely to be fatigued in the summer than an older player. This is particularly true for a league such as the NBA, which plays a large number of games (82) in the regular season. To take such aspects into account, I included the average number of years of experience in the NBA and age during the two seasons before the Games. Although these two variables are highly correlated, they do not represent the same thing. For example, a 22-year-old player might be in his first season or his fifth, depending on whether he went to the NBA directly from high school or stayed in college for four years before turning professional. To take into account possible non-linearities, I also included age and experience in squared terms. Additionally, I included a dummy variable for the five positions on the court: point guard, small guard, small forward, power forward, and centre. As the selection committee needs to create a balanced team, it must have a roughly fixed number of players in each position. To further balance physical characteristics, I also included height (in centimetres) and weight (in kilograms). I also included the average number of total games played – in the regular season and the playoffs – in the two years before the games as covariates. This variable allows the matching of selected players with controls that are similar in terms of various unobservables, such as physical condition and the player's history of injuries. Additionally, the inclusion of playoff games indirectly controls for the rest time that players had during the summer. The regressions also include the number of wins that each player's team had in the regular season. In this way, I was able to take into account the committee's possible preference for athletes from winning teams. I estimated the propensity score using a probit regression, including all the US athletes who played in at least one game in a season. A player might be a control in more than one Olympic Games, so the regressions employ standard errors clustered at the player level (Cameron and Miller, 2015). Selected players were included as treated only in the edition(s) in which they played, and are listed as controls in the other editions – if they met the criteria. For example, Michael Redd is included as treated in Bejing 2008 but as a control in 2004. All regressions include individual fixed effects.

The average marginal effects of the probit estimation can be seen in Column 1 of Table 2. To avoid having too many zero coefficients – and only in this table – all the controls have been divided by 100. The PER is by far the best predictor of selection for treatment. An increase by one unit of the PER leads to an increase of 1.4 percent in the probability of being selected. Trends in the PER also predict participation in the Games, although not as strongly as the baseline PER.[11] The number of games played is also important in the choice of Olympic athletes. In addition, taller players are more likely to be selected. Interestingly, age and experience are not significant, even though they have the expected sign. The number of wins of a team seems not to influence the decision to choose a player. In Column 2, I ran the same regression as in Column 1 but included two alternative performance measures: Box Plus/Minus (BPM) and Win Shares 48 (WS48).[12] Both measures synthesize the contribution of the player to the success of his team in terms of point differentials (BPM) and team wins in a season (WS48). In this way, I could stage a contest to identify the most important performance measure in the selection criteria. The results show that the PER is the only positive and significant measure of the three. The results in Columns 1 and 2 include all players, although, in the real world, the selection committee chooses from among a more restricted group of players. Are the criteria different if only plausible candidates are considered? To answer this question, I followed two different strategies. In the first, I exploited a change in the selection process that was made in 2008. From this edition onward, the 12 players selected for the Games were chosen from a pool of finalists. The size of the pool was different in each edition but was usually around 30–40 players (see the Online Appendix). Therefore, in Column 3, I only consider this latter group of players. The PER is still the most important variable considered by the Olympic committee, even among more homogeneous players. The trend in performance seems to lose importance. In Column 4, I restrict the analysis to the season before the Games, which allows me to increase the pool of potential controls. Finally, in Column 5, I consider the three years leading up to the Games.[13] Overall, Columns 2–5 confirm the results found in Column 1.

This work employs the propensity scores to calculate the kernel weights (Heckman et al., 1997), which are constructed in the following way:

$$\text{kernel weights}_i = \frac{K[(p_i - p_k)/h_n]}{\sum K[(p_i - p_k/h_n]}.$$

Here, p_i and p_k are the propensity scores for the treated and control units, K is the kernel function (the gaussian in this case), and h_n is the bandwidth,

[11] I excluded players who had trends greater or lower than 200 percent to avoid outliers. Nevertheless, no treated individuals were affected by the restrictions. The results are robust with other thresholds but also without restrictions.

[12] For more details, see the Basketball Reference website, About Box Plus/Minus (BPM), https://www.basketball-reference.com/about/bpm.html.

[13] Columns 2, 4, and 5 refer to the probit results for Columns 4, 5, and 6 of Table 5.

Table 2. Determinants of the participation into the Games: average marginal effects in a probit model

	Baseline (1)		All performances (2)		Original controls (3)		One year before (4)		Three years before (5)	
PER	1.135***	[0.113]	1.052***	[0.195]	1.890*	[0.977]	0.776***	[0.103]	1.211***	[0.125]
PER trend	0.043**	[0.022]	0.037*	[0.020]	−0.682	[2.726]				
Age	−1.212	[2.459]	−1.246	[2.519]	0.408	[2.001]	−0.551	[1.914]	−0.582	[2.689]
Games played	0.133***	[0.034]	0.131***	[0.031]	0.020	[0.028]	0.127***	[0.028]	0.097***	[0.037]
Experience	0.818	[0.658]	0.907	[0.672]	−0.180	[0.469]	0.919*	[0.555]	0.492	[0.758]
Point guard	0.048	[1.606]	−0.011	[1.585]	1.810	[1.546]	0.748	[1.196]	0.005	[1.673]
Small forward	1.971	[1.353]	1.979	[1.357]	−0.182	[1.170]	−0.320	[1.103]	1.822	[1.488]
Power forward	−2.087	[1.698]	−2.102	[1.678]	−4.451***	[1.289]	−3.574***	[1.297]	−2.924	[1.891]
Centre	−1.378	[2.251]	−0.976	[2.188]	−3.569**	[1.400]	−3.952**	[1.712]	−1.229	[2.384]
Height	0.107	[0.096]	0.109	[0.093]	0.171*	[0.092]	0.138*	[0.080]	0.143	[0.104]
Weight	−0.063	[0.055]	−0.070	[0.055]	0.016	[0.060]	0.007	[0.049]	−0.091	[0.060]
Experience squared	−0.039	[0.057]	−0.044	[0.058]	0.040	[0.038]	−0.030	[0.045]	−0.021	[0.063]
Age squared	0.013	[0.047]	0.014	[0.048]	−0.017	[0.039]	−0.001	[0.036]	0.002	[0.050]
Team wins	−0.032	[0.043]	−0.028	[0.044]	0.000	[0.029]	0.003	[0.028]	0.013	[0.049]
WS48			−0.081***	[0.013]						
BPM			0.127	[0.292]						
WS48 trend			−0.020	[0.014]						
BPM trend			−0.006	[0.006]						
PER trend (two years)									−0.006	[0.016]
N	1,778		1,752		244		2,176		1,561	

Notes: The dependent variable takes the value of one if the player went to the Olympic Games, and zero otherwise. All players with US passports in the season(s) before the summer of the Olympic Games are included. All regressions include the seven Olympic editions from 1992 until 2016. The control variables represent the pre-treatment characteristics most likely to have affected the selection – and participation – of a player into the Games. To ease the interpretation of the coefficients, all the controls have been divided by 100, except WS48 trend and BPM trend, which have been divided by 1,000 (these variables are defined in Section 2). Columns 1 and 2 consider the two seasons before the summer of the Games, but Column 2 also adds two additional measures of performance: WS48 and BPM. Column 3 considers only the players in the pool of finalists from the 2008 edition onward. Column 4 uses only the season before the Games, whereas Column 5 considers the percentage change in performance between three and one season before the Games. Standard errors, reported in brackets, are clustered at the individual level. ***, **, and * denote significance at the 1, 5, and 10 percent levels, respectively.

Table 3. Difference in means between treated and (weighted) controls

Weighted variable(s)	Mean control	Mean treated	Difference	t	$Pr(T > t)$
PER	21.29	21.52	0.233	0.40	0.691
PER trend	5.350	3.060	−2.290	0.45	0.653
Age	26.42	26.10	−0.327	0.67	0.502
Games played	77.79	80.22	2.438	1.11	0.268
Point guard	0.215	0.216	0.001	0.01	0.990
Small guard	0.215	0.196	−0.019	0.33	0.745
Small forward	0.229	0.270	0.041	0.64	0.520
Power forward	0.161	0.155	−0.005	0.11	0.911
Centre	0.200	0.176	−0.025	0.41	0.680
Experience	5.466	5.230	−0.236	0.52	0.604
Height	200.8	201.0	0.245	0.19	0.851
Weight	100.2	100.6	0.373	0.18	0.857
Age squared	39.65	691.5	−3.133	0.53	0.594
Experience squared	709.8	36.52	−18.29	0.70	0.484
Team wins	45.52	45.63	0.106	0.07	0.947

Notes: This table aims at showing covariate balance between the treated and control groups. Column 1 reports the mean for the treated, whereas Column 2 is for the (weighted) control group. Column 3 shows the difference between Columns 1 and 2. Column 4 reports the t-values of the t-test, whereas Column 5 reports the p-values.

which I set to 0.05.[14] I use the propensity scores calculated in Table 2 to calculate the weights. In Table 3, I perform a standard t-test of the difference in means between the covariates of treated and (weighted) controls before the Olympic Games. The results of this exercise show that the covariates are balanced between the two groups. Treated and controls are similar in terms of PER, which provides further evidence of the presence of a large overlapping area.

3. Main results and threats to identification

3.1. Econometric strategy and main results

In this section, I formally analyse the impact of going to the Olympic Games on the performance of selected players. Given that the treatment – the Olympic Games – took place between two seasons, I can identify a before and an after. Therefore, I use a difference-in-difference approach, employing the kernel weights for the controls in all periods. Treated individuals always have a weight of one. By combining these two methods – difference-in-difference and propensity score matching – I can take account of individual time-invariant unobserved heterogeneity and obtain

[14]In Table 3 of the Online Appendix, as a robustness, another kernel function, the Epanechnikov is considered.

a comparable control group (Smith and Todd, 2005). Other econometric approaches could have been used, such as the inverse probability weighting or the nearest-neighbour matching difference-in-difference. The results are robust to the use of alternative methods, as I will show. I calculate the average treatment effect for the treatment group, or ATT. Given that the interest is particularly in the impact for the selected players, and that there are many more control players than treated players, ATT is to be preferred.

Formally, the model tested is

$$\text{PER}_{i,t} = \alpha \text{Selected}_i + \beta \text{After}_t + \theta \text{Selected}_i * \text{After}_t + X_{i,t}\gamma + \varepsilon_{i,t}, \quad (1)$$

where i stands for the player and t for the season (one before and one after the Olympics in question). "Selected" is a binary variable equal to one if the player participated in the Games, and equal to zero otherwise. "After" is a dummy equal to one for the season after the Olympic Games. θ is the coefficient of interest, and $X_{i,t}$ is a vector of control variables. The analysis is restricted to those individuals in the common support (i.e., where the conditional distributions of X_i given treatment and controls overlap). Standard errors, clustered at the player level, are used in all specifications. I present the results in Table 4 by slowly adding control variables to assess the robustness of θ. Column 1 does not include any of the covariates in X. In Column 2, individual player fixed effects are added, to capture time-invariant characteristics. From Column 3 onward, all the remaining variables are progressively added. The last column is the preferred specification and, together with individual characteristics, includes fixed effects concerning the edition of the Olympic Games, the team, and the season.[15]

Participating in the Olympic Games has a strong positive and significant impact on the performance of treated individuals in the seasons following the Games. The coefficient is consistent across all the specifications. In Column 6, the ATT is 1.543, which represents an increase in performance of 7.1 percent. This result is statistically and economically significant.

3.2. Threats to identification

Although the results presented in the table are robust to the inclusion of several controls, it is worth considering possible threats to the identification strategy. Table 5 reports such exercises, which include – unless stated differently – the same controls as in Column 6 of Table 4. The first threat, common to all studies with this research design, relates to the parallel trend assumption. Although not directly testable, I provide evidence to support it. In Column 1 of Table 5, I run a difference-in-difference similar to Table 4, but setting the (placebo) treatment between two and one years

[15]Team fixed effects also take into account coach decisions.

Table 4. Baseline results

	PER (1)	PER (2)	PER (3)	PER (4)	PER (5)	PER (6)
Selected	0.625	0.012	−0.326	−0.194	−0.378	−0.419
	[0.548]	[0.566]	[0.595]	[0.518]	[0.530]	[0.529]
After	−1.333***	−1.333***	−1.074**	−1.120**	−1.301***	−0.695
	[0.403]	[0.481]	[0.547]	[0.507]	[0.499]	[0.554]
Selected × After	1.628***	1.628**	1.592**	1.441**	1.419***	1.543***
	[0.531]	[0.635]	[0.632]	[0.570]	[0.543]	[0.566]
Age			−1.442*	0.617	3.526*	3.171*
			[0.748]	[1.725]	[1.863]	[1.655]
Games played			0.031	0.027	0.021	0.023
			[0.027]	[0.030]	[0.034]	[0.034]
Experience			1.304*	0.921	0.687	0.897
			[0.757]	[1.134]	[1.282]	[1.247]
Height			−0.391	−0.835**	−1.185***	−1.055***
			[0.292]	[0.397]	[0.381]	[0.400]
Weight			−0.107*	−0.167***	−0.105	−0.083
			[0.057]	[0.057]	[0.070]	[0.071]
Age squared				−0.028	−0.082***	−0.080***
				[0.020]	[0.024]	[0.020]
Experience squared				−0.008	0.035	0.034
				[0.025]	[0.032]	[0.029]
Team wins				0.018	0.033	0.039
				[0.027]	[0.031]	[0.030]
Individual fixed effects	No	Yes	Yes	Yes	Yes	Yes
Position fixed effects	No	No	No	Yes	Yes	Yes
OG fixed effects	No	No	No	No	Yes	Yes
Team fixed effects	No	No	No	No	Yes	Yes
Season fixed effects	No	No	No	No	No	Yes
Observations	1,968	1,968	1,968	1,968	1,968	1,968
Adjusted R^2	0.037	0.710	0.727	0.746	0.765	0.771
Mean outcome at $t = 0$ (treated)	21.76	21.76	21.76	21.76	21.76	21.76
Effect relative to the mean	7.48%	7.48%	7.32%	6.62%	6.52%	7.09%

Notes: The table reports the results of the kernel matching difference-in-difference. These regressions aim at evaluating the impact of participating in the Olympic Games on the performance of the selected players the season following the Games. Each regression includes one observation for selected and control players for the seasons before and after the summer of the Olympic Games. All regressions refer to the seven Olympic editions from 1992 until 2016. The weights have been assigned based on the propensity score calculated in Column 1 of Table 2. The kernel function employed to calculate the weights is the gaussian one. In all the regressions, the dependent variable is the PER, which is a linear weight metric that summarizes the player's performance. Each column adds additional controls variables compared with the previous column. Standard errors, reported in brackets, are clustered at the individual level. ***, **, and * denote significance at the 1, 5, and 10 percent levels, respectively.

before the Games. For the 1992 Olympic Games, this means comparing the seasons 1990–91 and 1991–92. In Column 2, a similar approach is used but considering the placebo treatment between one and two seasons after. Continuing with the 1992 example, I compare the seasons 1992–93 and 1993–94. In Column 3, I consider the same placebo treatment as in Column 1, but compare it with controls in Olympic years. Considering the 1992 edition, this means comparing the treated players in seasons 1990–91 and 1991–92 with the controls in seasons 1991–92 and 1992–93. Even though the timing between the (placebo) treated and non-treated is different, I do not expect the former to show a significant change in performance. I do not find any effect in all these exercises, which supports the view that there were no different trends between treated and controls in a period other than the actual treatment. Continuing the investigation, what if the PER is not the only performance variable that the selection committee takes into account? In Column 4, I include the two measures presented earlier – BPM and WS48 – along with their trends: the coefficient θ is still positive and significant. If anything, these measures are slightly higher than the baseline in Column 6 of Table 4. Given that I condition on a broader definition of performance, it is reassuring that the impact of going to the Olympic Games has a relevant economic significance. Furthermore, the observed results might be driven by, or be at least sensitive to, the choice of the covariates included for the computation of the propensity score. In the baseline analysis, I use the averages for the two years before the Games, which leads to the exclusion of those players with only one season of NBA experience from the control group. To check whether such exclusion is affecting the results, in Column 5, I use only the covariates for the season before the Games to calculate the propensity score. This means including all the variables presented in Section 2.2, except the trend because it requires two seasons to be calculated. The interaction coefficient is still positive and significant, while the number of observations has increased. The opposite argument can also be made: the selection committee might look at the performance progression not only in the two seasons before the Games but in the three seasons before. Thus, I calculate the trend between three and one seasons before the Games while keeping the average of the other covariates for two seasons.[16] The coefficient, presented in Column 6, is still positive and strongly significant, although slightly smaller than the one in the last column of Table 4. In Column 7, I consider an inverse probability weighting difference-in-difference technique. The weights are

[16]In Table 2 of the Online Appendix, I report two additional robustness exercises. In Column 7, I include two trends, between the seasons $t-3$ and $t-2$, and $t-2$ and $t-1$ (where t is the summer of the Olympics). In Column 8, I consider the average between the two trends. The results are similar to those in the main body of this work.

Table 5. Threats to identification

	Placebo before (1)	Placebo after (2)	Placebo mixed (3)	Different measures (4)	One year (5)	Three years (6)	IPW (7)	Injured (8)	No 2004 (9)
Selected × After	0.513	−0.591	−0.018	2.138***	1.472***	1.615***	1.705**	−0.173	1.950***
	[0.454]	[0.460]	[1.071]	[0.661]	[0.419]	[0.572]	[0.672]	[0.991]	[0.671]
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,968	2,547	1,968	1,918	2,528	1,546	1,968	632	1,324
Adjusted R^2	0.813	0.809	0.778	0.711	0.810	0.723	0.807	0.785	0.703

Notes: This table reports various exercises to assess the robustness of the baseline regression to potential threats to the identification strategy. The dependent variable in all the nine regressions is the PER, which is a linear weight metric that summarizes the player's performance. Column 1 sets the placebo treatment between two seasons and one season before the Olympic Games. Column 2 sets the placebo treatment between the first and second seasons after the Games. Column 3 sets the placebo treatment as in Column 1 but the controls refer to the true Olympic years. In Column 4, I condition also on WS48 and BPM, and the weights have been calculated from the results in Column 2 of Table 2. Column 5 employs the weights calculated as for Column 4 in Table 2. In Column 6, the weights are calculated based on Column 5 in Table 2. Column 7 uses, as weights for the difference-in-difference exercise, the inverse of the probability of being selected. Column 8 considers as treated those players who were selected to go to the Games but could not participate because of injuries. Finally, Column 9 excludes the selected players from the 2004 Games. All the regressions include the set of control variables employed in Column 6 of Table 4. These are age, age squared, experience, experience squared, games played, height, weight, position on the court, and team wins. The regressions also include individual, team, Olympic Game edition, and season fixed effects. Standard errors, reported in brackets, are clustered at the individual level. ***, **, and * denote significance at the 1, 5, and 10 percent levels, respectively.

the inverse probability of being selected in the treatment. The main result is maintained, and the coefficient is somewhat bigger.

Finally, another problem is that Olympic players might have increased their performance irrespective of their participation in the Games. To rule out this possibility, I run a falsification test and consider as treated those players who were named to be part of the Olympic men's basketball team but could not participate because of injuries. Finding that such players increased their performance after the Games, even though they did not participate, would raise queries regarding the causal role of the Olympics in explaining the results.[17] In total, 15 players could not participate in the Games because of injuries.[18] The results of this exercise, in Column 8, show that injured players did not have a statistically different performance from the controls. Furthermore, in the 2004 Olympic Games, many of the very

[17]For the editions from 1992 until 2004, I consider the players who were named among the 12 but who were replaced. For example, in 1996 Gary Payton replaced Glenn Robinson with an Achilles' tendon injury. For the editions 2008, 2012, and 2016, I consider the roster finalists who had to withdraw because of injuries. This is the case of Anthony Davis in 2016. I thank Craig Miller for providing me with the data for the editions between 1992 until 2004.

[18]By construction, all the treated individuals played the season after the Games, where they played an average number of 62 games, which is higher than for non-Olympic athletes. This means that injuries did not affect the performance of Olympic players.

best players selected out. Most of them chose not to play citing personal reasons, and their absence could affect the results by decreasing the average quality of the Olympic team and, thus, limiting the role of peer effects.[19] To assess whether this is the case, in Column 9, I ran the baseline model but excluded the 2004 Olympic players. The coefficient is now 1.905, which is higher than the one found in Column 6 of Table 4. If anything, the inclusion of 2004 lowers the role of peer effects. The next step is to provide some additional results to contextualize the findings.

4. Further results, dynamics, and heterogeneity

This section presents the additional results along with some evidence on the dynamics and heterogeneity of treatment, which can be found in Table 6. The first set of regressions studies whether participating in the Olympic Games affects other (advanced) measures of performance. Column 1 considers BPM, whereas Column 2 considers WS48. The covariates employed to calculate the propensity score are similar to the ones used in Table 2, except for the performance measure and its related trend. The coefficients are still positive and strongly significant. For example, those who participated in the Olympic Games increased their BPM and WS48 by 13.5 percent and 13.6 percent, respectively, which is higher than the effect found for the PER. These results show that going to the Olympics also affects other dimensions of a player's contribution to the success of the team. In Column 6 of Table 3 of the Online Appendix, I conduct a further exercise using another performance measure, Win Shares (WS), which reveals a similar result.

The following exercises deal with the dynamics of the treatment. In Column 3, I consider the two seasons before and after the Olympic Games, interacting the treatment with a dummy for each of these seasons. The results, also presented in Figure 1, reveal that there is a decrease in performance in the second season after the Games. This exercise also shows that two seasons before the Games, there were no statistically significant differences between the treated and controls, which further confirms the presence of parallel trends. In Column 4, I study the role of potential heterogeneous effects for each Olympic Games and interact "Selected × After" with a dummy for each of the seven editions, with the 1992 edition as the excluded category. The analysis reveals that the coefficients are statistically significant for the 1996, 2008, and 2012 editions, and not for the others. In Column 5, I replicate the model in Column 4 but exclude the 2004 Games for the reasons explained in the previous section. The results

[19] Jason Kidd was the only top player who could not participate because of injuries.

are similar to those in Column 4. In Column 6, I consider only the controls coming from the pool system explained in Section 2.2. The results confirm that Olympic players show an increase in performance.[20]

NBA players typically spend most of their offseason resting, which they obviously cannot do when participating in the Olympic Games. As such, the increase in performance could be potentially explained by the heterogeneous level of practice between the selected and controls. To test whether practice during the summer plays a role, I focus exclusively on a restricted group of control players who had been engaging in summer official tournaments. More specifically, I gathered information on all the NBA players who, during the summers of the Games, participated in any of the NBA summer leagues. There are three main such competitions: the Las Vegas, Orlando Pro, and Salt Like City summer leagues.[21] The summer leagues are intended to feature try-out players who could fill some spots in the upcoming NBA regular season line-ups.[22] Often these tournaments include experienced NBA players who want to keep in shape during the offseason. In Column 7, I compare selected players only with NBA players who (a) participated in one of the three summer leagues and (b) played in an NBA team in the season before and after the Games.[23] Unfortunately, summer leagues have existed only since 2002, which means that I can consider only the Olympic Games editions from 2004 onward. The result shows a positive and significant effect, with a similar magnitude to the baseline. Furthermore, I test whether the players with relatively lower skills benefit more than those with higher skills from participating in the Olympic Games. Even though selected players are all very talented, the impact on Michael Jordan might be different from the impact on the less talented Chris Mullin, both of whom were in the 1992 Dream Team. Thus, I divide players into two groups: below and above the median skills' level. Then I run a triple difference-in-difference for those below the median, reported in Column 8. I did not find any effect, which indicates the homogeneity of the treatment effects across skill abilities.[24]

[20] In Table 2 of the Online Appendix, I run two additional exercises. In Column 1, I include the players who participated in multiple editions of the Olympic Games only once, in his first appearance. In Column 2, I consider only those players that have been selected at some point in their career. In both cases, I find a significant effect of participating in the Olympic Games.

[21] The Las Vegas league is by far the most famous and respected of the three.

[22] First-year players – rookies – usually participate in such tournaments, even if they were picked high in the draft.

[23] Given the limited number of players, I do not identify controls through a matching technique and simply assign a weight of one to all controls. This exercise should be considered only as indicative evidence because I cannot rely on the balancing of covariates as in the baseline.

[24] In Columns 7 and 8 in Table 3 of the Online Appendix, I provide further evidence for this, employing a quantile regression.

Table 6. Robustness checks

	BPM	WS48	Dynamics			Original	Summer	Skills
			Year	OG edition	OG edition (no 2004)	controls	league	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Selected × After	0.557*	0.024***		−0.624	−0.806	1.313**	1.625*	1.578***
	[0.316]	[0.008]		[0.820]	[0.913]	[0.529]	[0.976]	[0.566]
Two years before			−0.935*					
			[0.505]					
One year after			1.145***					
			[0.402]					
Two years after			0.378					
			[0.433]					
Selected × After × 1996				6.340**	5.705**			
				[2.657]	[2.283]			
Selected × After × 2000				1.081	1.298			
				[1.027]	[1.165]			
Selected × After × 2004				0.975				
				[1.203]				
Selected × After × 2008				2.675*	2.763*			
				[1.400]	[1.447]			
Selected × After × 2012				2.737*	3.191**			
				[1.410]	[1.582]			
Selected × After × 2016				0.956	0.766			
				[1.230]	[1.328]			
Selected × After × Below median								−0.297
								[1.236]

Table 6. Continued

	BPM	WS48	Dynamics			Original	Summer	Skills
			Year	OG edition	OG edition (no 2004)	controls	league	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2,274	2,768	3,881	1,968	1,386	248	552	1,968
Adjusted R^2	0.784	0.727	0.738	0.789	0.797	0.472	0.779	0.771

Notes: This table reports various robustness exercises. Columns 1 and 2 replicate the baseline analysis presented in Column 6 of Table 4 with BPM and WS48 – rather than the PER – as the dependent variable. The propensity scores have been calculated considering such performance measures. Column 3 extends the analysis to two seasons before and two after the Games and interacts the treatment variable with season dummies. Column 4 shows the heterogeneity of response to treatment depending on the edition of the Olympic Games. Column 5 replicates the exercise in Column 4 but excluding the selected players in 2004. Column 6 considers only the players that were in the pool of candidates in the two to three year programs after 2008. Column 7 uses, as controls, only the players that participated in one of the three NBA summer leagues for the Games from the 2004 edition onward. Column 8 reports the triple interaction for the players below the median level of skills. The dependent variable in Column 1 is BPM, which is a Plus/Minus statistic, whereas WS48 (in Column 2) represents the player's contribution to the wins of his team in each season. The dependent variable for Columns 3–8 is the PER, which is a linear weight metric that summarizes the player's performance. All the regressions include all the control variables employed in Column 6 of Table 4. These are age, age squared, experience, experience squared, games played, height, weight, position on the court, and team wins. The regressions also include individual, team, Olympic Games edition, and season fixed effects. Standard errors, reported in brackets, are clustered at the individual level. ***, **, and * denote significance at the 1, 5, and 10 percent levels, respectively.

Figure 1. Dynamics of PER two seasons before and after the Olympic Games

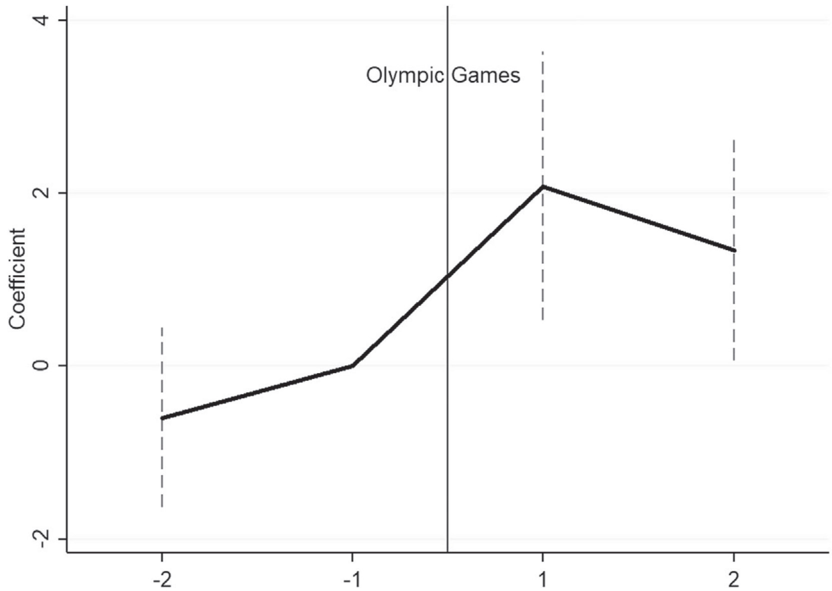

Notes: The figure presents the coefficients of the dynamic specification in Column 3 of Table 6. It shows the interaction between the dummy "Selected" with the two seasons before and after the summer of the Olympic Games. The omitted season is *-1*. The dependent variable is the PER, which is a linear weight metric that summarizes the player's performance. The regressions include age, age squared, experience, experience squared, games played, height, weight, position on the court, and team wins. The regressions also include individual, team, Olympic Games edition, and season fixed effects. The dotted vertical line represents the 95 percent confidence intervals.

5. Channels

The previous sections have demonstrated that the Olympic players studied increased their performance compared with those who did not participate in the Olympic Games. The identification exercises provided findings that can be taken as causal. The next question to consider is: what explains these results? What is so special about the Games that they make performances improve? I argue that peer effects are the most likely factors to explain the increase in performance. During the NBA tournament, the Olympic athletes played alongside teammates who were, on average, of much lower quality. The average PER of the Olympic players' teammates in their NBA team the season before the Games was 13.02. However, the average PER of the Olympic players was 21.75. This represents a positive shock, in terms of peer quality, of about 67 percent, which can help to explain why the performance of selected players increased after each Games. The positive shock was not homogeneous among the selected players. Those who had relatively lower-quality teammates during the regular NBA seasons were more likely to receive a greater peer shock compared with those who played in higher-quality teams. For example, in the 2015–16 season, Harrison Barnes, who participated in Rio 2016, was playing for the Golden State Warriors alongside other top players, such as Stephen Curry, Klay Thomson, and Diamond Green. Going to the Olympics did not, therefore, represent a major increase in peer quality for him. However, Patrick Ewing's teammates in 1991–92 were much weaker than the other members of the 1992 Olympic Team. It is therefore the size of this peer shock that explains the increase in performance, rather than the quality of an individual's teammates per se.

To test these claims, I constructed the variable "Peer shock" for the 79 Olympic players, which I calculated by subtracting the peer quality of the players in the Olympic team from that of the NBA team in the season before the Games. In Column 1 of Table 7, I regress the change in performance between the seasons before and after the games against the Peer shock, but only for the 79 Olympic players. The results show that the performance improvement is positively associated with the Peer shock.[25] An increase of one unit of the latter variable causes an increase of 0.568 in the PER, a sizeable effect. This effect is similar to the one found by Mas and Moretti (2009), who revealed that a 10 percent increase in co-worker productivity leads to a 1.5 percent increase in one's productivity. In my case, the increase is 2.27 percent for 10 percent. In Column 2, I distinguish between the Peer shock for those below and above the median, and I assess its interaction with the time and treatment variables. I then run a kernel-weighted difference-in-difference regression using all the control players. The results confirm that those with Peer shock above the median performed much better than those below.[26] In Column 3, I test whether the quality of the Olympic teammates, rather than the shock in peer quality, is driving the results. I run a regression similar to that in Column 1 but using the average PER of teammates the season before the Games. As we can see, the peer quality of Olympic teammates does not explain on its own why Olympic players increase their performance. Furthermore, between 1992 and 2016, 78 non-US NBA players competed at the Games with their national teams and also satisfied the criteria needed to be part of the analysis. Compared with the US players, the non-US athletes played in weaker national teams. Did these non-US Olympic players also experience a boost in performance in the NBA season after the Games? In Column 4, I run a difference-

[25]I am aware that the typical peer effect model is the linear-in-means social interaction model. This involves regressing the individual outcome on the average outcome of peers plus a set of individual explanatory variables, including past individual outcome levels (Sacerdote, 2011). I cannot use this strategy for the Olympic Games because of the unavailability of advanced statistics and the low number of matches – generally non-competitive – that are played during the Games.

[26]The difference between these two coefficients is statistically different from zero.

in-difference regression using as treated the non-US NBA players who played in the Summer Olympic Games. It shows that these players did not increase their performance, which supports the role of the peer effect.

Despite this evidence, there might be other factors that explain the improved performance of US Olympic athletes. The Olympic team included not only elite basketball players, but also elite coaches. All five of the Olympic coaches during this period were Hall of Famers, and each had won at least one NCAA or NBA championship. These coaches were much better coaches than those in the NBA teams. In a sense, the selected players experienced a positive coaching shock, which might have affected their performance (Lazear et al., 2015). To test such a claim, I employed a measure of shock that reflects the difference in performance between the Olympic coach and the NBA coach for the selected players. This measure is based on the percentage of wins, given by the total number of wins divided by the total number of games coached.[27] I considered both regular seasons and playoff games.[28] For each manager, I considered their record up to the summer of the Olympic Games. On average, the Olympic coaches had a win percentage 10.96 percent higher than their non-Olympic counterparts. I employ this measure of coaching shock similarly to Column 1 of Table 7. The results in Column 5 reveal that the increase in coach quality is not a channel through which selected players increase their performance. I also explored whether the level of competition from the Olympic opponents played a role. To do so, for each edition of the Games, I calculated the number of NBA players that the US team faced during the tournament. Then I performed a standard difference-in-difference analysis, allowing this variable to interact with Selected × After. The results are shown in Column 6 and reveal an absence of any effect on performance. In Column 7, I consider the average point differentials between Team USA and their opponents. This factor also does not seem to explain the results.

Columns 1–7 suggest that the shock in peer quality between the NBA and Olympic teams is the most convincing explanation for the increase in performance after the Games. In the last two columns of Table 7, I explore whether learning effects are a possible channel. Following the exercises used in the literature on peer effects, such as Cornelissen et al. (2017) and

[27] I considered the records in both the NBA and NCAA. Several NBA coaches have had long and successful careers as NCAA coaches. This is the case with Larry Brown, Billy Donovan, and Jim Lynam, while Mike Krzyzewski never coached an NBA team, but is considered to be one of the most successful coaches in the history of the game. However, I restricted the analysis to NCAA coaches in the first division. If a team had more than one manager in the season before the Olympic Games, I weighted the percentage according to the number of games coached by each coach. Finally, the same person could be included as both an Olympic and a non-Olympic coach, as is the case with Larry Brown.

[28] For the NCAA, I consider as playoff games those played in the NCAA tournament bracket.

Table 7. Channels

	Peer quality change I (1)	Peer quality change I (2)	Average OG quality (3)	Non-US players (4)	Coach quality change (5)	Average NBA opponents (6)	Points differentials (7)	Old vs young (8)	Experienced players (9)
Peer shock	0.568**								
	[0.247]								
Above × After		1.727***							
		[0.596]							
Below × After		0.408							
		[0.711]							
Average quality			−0.679						
			[0.898]						
Selected non-US × After				−0.310					
				[0.446]					
Coaching shock – % wins					−0.007				
					[0.034]				
Selected × After × Average NBA opponents						−0.025			
						[0.298]			
Selected × After × Points differential							0.011		
							[0.036]		
Selected × After × Old								1.347	
								[1.133]	
Selected × After × Experience									−0.275
									[1.130]

Table 7. Continued

	Peer quality change I (1)	Peer quality change I (2)	Average OG quality (3)	Non-US players (4)	Coach quality change (5)	Average NBA opponents (6)	Points differentials (7)	Old vs young (8)	Experienced players (9)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	79	1,968	79	3,993	79	1,968	1,968	1,968	1,968
Adjusted R^2	0.239	0.773	0.179	0.696	0.173	0.771	0.771	0.779	0.778

Notes: This table reports various exercises that explore possible channels for the increase in the performance of Olympic athletes. Columns 1–4 refer to the peer effect explanations. Column 1 regresses the changes in the PER between the seasons before and after the Games on Peer shock (i.e., the difference in the average quality of teammates between the NBA and the Olympic teams). It does so only for the Olympic players. Column 2 considers a regression similar to Column 6 in Table 4 but separating the treated between those with Peer shock above and below the median. Column 3 regresses the change in the PER for the Olympic athletes on the average quality, in terms of the PER, of the Olympic team. The following column evaluates the impact of going to the Olympic Games on the non-US players. Column 5 explores the role of coaching shock, by measuring the difference in wins percentage between the coaches of the Olympic and NBA teams. Columns 6 and 7 explore alternative specifications. The former interacts Selected × After with the average number of NBA players that the US team faced in each Games. Column 7 interacts Selected × After with the average point differential between the US team versus all its opponents in the Olympic tournament. A lower average differential is a signal of more competition and better overall quality. The following two columns explore whether the effect of participating in the Olympic Games is heterogeneous depending on age and experience. The coefficients are from a triple difference-in-difference model. "Old" is a binary variable equal to one if the player was older than 27 years old before the Games. "Experienced players" is a binary variable equal to one if the athlete had six, or more, years of experience in the NBA. In all the columns, the dependent variable is the PER, which is a linear weight metric that summarizes the player's performance. All the regressions, except for Columns 1, 3, and 5, include all the control variables employed in Column 6 of Table 4. These are age, age squared, experience, experience squared, games played, height, weight, position on the court, and team wins. The regressions also include individual, team (except in Columns 1, 3, and 5), Olympic Games edition, and season fixed effects. Standard errors, reported in brackets, are clustered at the individual level. ***, **, and * denote significance at the 1, 5, and 10 percent levels, respectively.

Brune et al. (2022), I assess the heterogeneity across age and experience. More specifically, in Column 8, I define the binary variable “Old”, which takes a value of 1 if the player was at least 27 years old during the season before an Olympic Games.[29] In Column 9, I define the dummy “Experienced players”, which takes a value of 1 if the athlete had at least six years of experience in the league before a Games. Given that such a dummy defines a further characteristic, I consider a triple difference-in-difference. In both specifications, I fail to find any significant effects across such dimensions. These results suggest that knowledge transfer has a limited effect, similar to the findings of Jiang (2020) for swimmers. In Column 5 of Table 2 in the Online Appendix, I consider only the players who were over 30 in the year of the Olympic Games. The results show that this effect still holds. In Column 6 of Table 2 in the Online Appendix, I consider a different threshold of experience (i.e., 5 instead of 6). Again, I do not find any significant difference.

In the following section, I determine whether there were positive externalities for the NBA teams in the season after the Olympic Games.

6. Trickle-down effects

In this section, I assess whether the benefits of going to the Olympic Games extend from the Olympic athletes to their teammates in their original NBA teams. Are there positive trickle-down effects for lower-skilled players? The literature has shown mixed evidence on the impact of star workers on lower-skilled colleagues (Agrawal et al., 2017; Serafinelli, 2019). To answer this question, I define all the players in a team with an Olympic athlete as treated and those without as controls. Next, I run a standard difference-in-difference regression comparing the seasons before and after for these two groups, excluding from the sample the Olympic athletes. Results are reported in Table 8.

In Column 1, I consider all the players, irrespective of whether they changed teams between the season before and after. In Column 2, I restrict the sample only to the athletes in the same team between these two seasons. I do not find evidence of any effect in either specification. Next, I check whether the effect depends on the players’ skills. What to expect is not clear *a priori*: it might be that the players with the lowest skills are those most likely to be affected. To check this, I divide the players into five skills quintiles based on the PER in the season before the Games. I then run the same regression for each of these quintiles. Such exercises provide some robustness in the absence of positive spillover effects for lower-skilled teammates. In Column 8, I run a standard linear-in-means regression

[29] Results are consistent with other age thresholds, such as 27 and 26.

Table 8. Trickle-down effects

	All players (1)	No change team (2)	Q1 (3)	Q2 (4)	Q3 (5)	Q4 (6)	Q5 (7)	Peer quality (8)
Team with OG × After	0.237	0.187	0.585	−0.435	0.334	−0.322	−0.086	
	[0.226]	[0.280]	[0.804]	[0.443]	[0.418]	[0.376]	[0.537]	
Peer average quality								−0.168
								[0.108]
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	4,887	2,592	985	978	976	992	956	2,443
Adjusted R^2	0.556	0.620	0.242	0.233	0.224	0.223	0.400	0.624

Notes: This table reports various exercises to explore whether participating in the Olympic Games had positive trickle-down effects for Olympic players' teammates in the original NBA teams. Column 1 considers as treated all the players in the same team as the Olympic athletes and controls the others. Olympic players are excluded from the analysis. Column 2 restricts the sample to those players who did not change team between the season before and after the Games. Columns 3–7 separate players into five quintiles of PER and run five separate regressions. Column 8 reports the exercise with a naïve linear-in-mean model. In all regressions, the dependent variable is the PER, which is a linear weight metric that summarizes the player's performance. All the regressions include all the control variables included in Column 6 of Table 4. These are age, age squared, experience, experience squared, games played, height, weight, position on the court, and team wins. The regressions also include individual, team (except in Column 8), Olympic Games edition, and season fixed effects. Standard errors, reported in brackets, are clustered at the individual level. ***, **, and * denote significance at the 1, 5, and 10 percent levels, respectively.

(Sacerdote, 2011). The results show that the average performance of the team does not predict players' performance. I take this coefficient but urge caution, as I am fully aware that peer quality might suffer from endogeneity and reflection (Manski, 1993; Angrist, 2014).

The results point to the absence of trickle-down effects for the teammates of the Olympic players in the season after the Games. How can we explain such findings? Compared with the existing literature, this setting is unique. The superstar workers do not move to a different firm but experience a boost in productivity due to their participation in the Games. Their co-workers experience an indirect shock, which might not substantially affect their performance. Combining these findings with those in the previous section, it seems that the benefits of participating in the Olympic Games are private – confined to the Olympic players – and not public (i.e., they do not trickle down to other workers).

7. Conclusions

In this paper, I aim to assess whether peer interactions between superstar workers lead to an increase in performance. As a source of peer interaction, I considered the participation of elite NBA US players in the Olympic Games, which took place between two NBA seasons. I then evaluated the change in the performance of these players before and after the

Games. Using detailed information about individual advanced performance statistics, I found a sizeable increase in the performance for the selected players. Olympic players improved their performance by 7.1 percent compared with the control group. These findings are robust to different performance measures and control groups.

Once it was established that the results could be seen as causal, I started to explore potential channels, finding evidence that peer effects are the factor that is most likely to explain the results. Olympic players experience a positive shock in peer quality by going to the Games. I also explored alternative channels, such as the shock in the level of coaching, but could not identify a more convincing explanation than peer effects. Additionally, I assessed whether the (lower-quality) players who only play in the NBA benefited from playing alongside better Olympic athletes the season after the Games. The results show that there were no trickle-down effects.

In this work, I make a relevant contribution to the existing literature (Falk and Ichino, 2006; Gould and Winter, 2009; Guryan et al., 2009; Mas and Moretti, 2009; Waldinger, 2011; Serafinelli, 2019). I provide clear evidence of the benefits of interaction between highly skilled workers. I do so by employing individual performance statistics, which allow me to better evaluate the impact of the workers (players) on the success of their firms (teams). To my knowledge, this is the first time that such performance measures have been used to evaluate productivity in such a context. These results have some relevant implications. If firms want to increase the overall performance of their labour force, they should encourage the most talented employees to collaborate equally with less talented workers for some time. However, the firms must be aware that the benefits of such initiatives might not necessarily trickle down to other employees. If this is the case, then firms could think about co-payment methods together with the star employees involved. Finally, such an exercise could be replicated in many other labour economics contexts. For example, it could be used to evaluate changes in the productivity of the best academics when they collaborate with equally talented colleagues at other universities for extended research periods. This exercise could also be used to test professionals who regularly participate in government-organized task forces within their area of expertise.

Supporting information

Additional supporting information can be found online in the supporting information section at the end of the article.

Online appendix
Replication files

References

Agrawal, A., McHale, J., and Oettl, A. (2017), How stars matter: recruiting and peer effects in evolutionary biology, *Research Policy 46*, 853–867.

Angrist, J. D. (2014), The perils of peer effects, *Labour Economics 30*, 98–108.

Azoulay, P., Graff Zivin, J. S., and Wang, J. (2010), Superstar extinction, *Quarterly Journal of Economics 125*, 549–589.

Babcock, P., Bedard, K., Charness, G., Hartman, J., and Royer, H. (2015), Letting down the team? Social effects of team incentives, *Journal of the European Economic Association 13*, 841–870.

Bandiera, O., Barankay, I., and Rasul, I. (2010), Social incentives in the workplace, *Review of Economic Studies 77*, 417–458.

Becker, G. S. (2009), *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education*, University of Chicago Press, Chicago, IL.

Brown, J. (2011), Quitters never win: the (adverse) incentive effects of competing with superstars, *Journal of Political Economy 119*, 982–1013.

Brune, L., Chyn, E., and Kerwin, J. (2022), Peers and motivation at work: evidence from a firm experiment in Malawi, *Journal of Human Resources 57*, 1147–1177.

Cairney, J., Joshi, D., Li, Y.-C., and Kwan, M. Y. (2015), The impact of the Olympics on regular season team performance in the National Hockey League, *Journal of Athletic Enhancement 4* (6).

Cameron, A. C. and Miller, D. L. Miller (2015), A practitioner's guide to cluster-robust inference, *Journal of Human Resources 50*, 317–372.

Cornelissen, T., Dustmann, C., and Schönberg, U. (2017), Peer effects in the workplace, *American Economic Review 107* (2), 425–456.

De Grip, A. and Sauermann, J. (2012), The effects of training on own and co-worker productivity: evidence from a field experiment, *Economic Journal 122*, 376–399.

Depken, C. A. and Haglund, L. E. (2011), Peer effects in team sports: empirical evidence from NCAA relay teams, *Journal of Sports Economics 12*, 3–19.

Emerson, J. and Hill, B. (2018), Peer effects in marathon racing: the role of pace setters, *Labour Economics 52*, 74–82.

Falk, A. and Ichino, A. (2006), Clean evidence on peer effects, *Journal of Labor Economics 24*, 39–57.

Gould, E. D. and Winter, E. (2009), Interactions between workers and the technology of production: evidence from professional baseball, *Review of Economics and Statistics 91*, 188–200.

Guryan, J., Kroft, K., and Notowidigdo, M. J. (2009), Peer effects in the workplace: evidence from random groupings in professional golf tournaments, *American Economic Journal: Applied Economics 1*, 34–68.

Hamilton, B. H., Nickerson, J. A., and Owan, H. (2003), Team incentives and worker heterogeneity: an empirical analysis of the impact of teams on productivity and participation, *Journal of Political Economy 111*, 465–497.

Heckman, J. J., Ichimura, H., and Todd, P. E. (1997), Matching as an econometric evaluation estimator: evidence from evaluating a job training programme, *Review of Economic Studies 64*, 605–654.

Hickman, D. C. and Metz, N. E. (2018), Peer effects in a competitive environment: evidence from the PGA Tour, *Economic Inquiry 56*, 208–225.

Hollinger, J. and Hollinger, J. (2005), *Pro Basketball Forecast, 2005–06*, Potomac Books, Dulles, VA.

Ichniowski, C. and Preston, A. (2014), Do star performers produce more stars? Peer effects and learning in elite teams, National Bureau of Economic Research (NBER), Working Paper 20478.

Jackson, C. K. and Bruegmann, E. (2009), Teaching students and teaching each other: the importance of peer learning for teachers, *American Economic Journal: Applied Economics 1*, 85–108.

Jiang, L. (2020), Splash with a teammate: peer effects in high-stakes tournaments, *Journal of Economic Behavior & Organization 171*, 165–188.

Kahn, L. M. (2000), The sports business as a labor market laboratory, *Journal of Economic Perspectives 14* (3), 75–94.

Kaur, S., Kremer, M., and Mullainathan, S. (2010), Self-control and the development of work arrangements, *American Economic Review 100* (2), 624–628.

Kubatko, J., Oliver, D., Pelton, K., and Rosenbaum, D. T. (2007), A starting point for analyzing basketball statistics, *Journal of Quantitative Analysis in Sports 3* (3).

Kuehn, J. (2017), Accounting for complementary skill sets: evaluating individual marginal value to a team in the National Basketball Association, *Economic Inquiry 55*, 1556–1578.

Lazear, E. P., Shaw, K. L., and Stanton, C. T. (2015), The value of bosses, *Journal of Labor Economics 33*, 823–861.

Manski, C. F. (1993), Identification of endogenous social effects: the reflection problem, *Review of Economic Studies 60*, 531–542.

Mas, A. and Moretti, E. (2009), Peers at work, *American Economic Review 99* (1), 112–145.

Oliver, D. (2004), *Basketball on Paper: Rules and Tools for Performance Analysis*, Potomac Books, Dulles, VA.

Rosen, S. and Sanderson, A. (2001), Labour markets in professional sports, *Economic Journal 111*, 47–68.

Sacerdote, B. (2011), Peer effects in education: how might they work, how big are they and how much do we know thus far? in E. Hanushek, S. Machin, and L. Woessmann (eds), *Handbook of the Economics of Education*, Vol. 3, Chapter 4, Elsevier, Amsterdam, 249–277.

Serafinelli, M. (2019), “Good” firms, worker flows, and local productivity, *Journal of Labor Economics 37*, 747–792.

Smith, J. A. and Todd, P. E. (2005), Does matching overcome LaLonde’s critique of nonexperimental estimators? *Journal of Econometrics 125*, 305–353.

Waldinger, F. (2011), Peer effects in science: evidence from the dismissal of scientists in Nazi Germany, *Review of Economic Studies 79*, 838–861.

Yamane, S. and Hayashi, R. (2015), Peer effects among swimmers, *Scandinavian Journal of Economics 117*, 1230–1255.

First version submitted December 2020;
final version received January 2022.