

NBER WORKING PAPER SERIES

PEER EFFECTS IN THE WORKPLACE:
EVIDENCE FROM RANDOM GROUPINGS IN PROFESSIONAL GOLF TOURNAMENTS

Jonathan Guryan
Kory Kroft
Matt Notowidigdo

Working Paper 13422
<http://www.nber.org/papers/w13422>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2007

E-mails: jguryan@chicagogsb.edu, kroft@econ.berkeley.edu, noto@mit.edu. The authors thank Daron Acemoglu, David Autor, Oriana Bandiera, Marianne Bertrand, David Card, Kerwin Charles, Ken Chay, Stefano DellaVigna, Amy Finkelstein, Alex Mas, Sean May, Steve Pischke, Imran Rasul, Jesse Rothstein and Emmanuel Saez for helpful conversations regarding this paper. The authors also thank Phil Wengerd for outstanding research assistance. Guryan thanks the University of Chicago Graduate School of Business for funding support. Kroft thanks the Center for Labor Economics and The Institute of Business and Economic Research at Berkeley for funding support. Notowidigdo thanks the MIT Department of Economics for funding support. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research. This research was funded in part by the George J. Stigler Center for the Study of the Economy and the State at the University of Chicago Graduate School of Business.

© 2007 by Jonathan Guryan, Kory Kroft, and Matt Notowidigdo. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Peer Effects in the Workplace: Evidence from Random Groupings in Professional Golf Tournaments
Jonathan Guryan, Kory Kroft, and Matt Notowidigdo
NBER Working Paper No. 13422
September 2007
JEL No. J01,J24,J3,J44

ABSTRACT

This paper uses the random assignment of playing partners in professional golf tournaments to test for peer effects in the workplace. We find no evidence that the ability of playing partners affects the performance of professional golfers, contrary to recent evidence on peer effects in the workplace from laboratory experiments, grocery scanners, and soft-fruit pickers. In our preferred specification, we can rule out peer effects larger than 0.045 strokes for a one stroke increase in playing partners' ability, and the point estimates are small and actually negative. We offer several explanations for our contrasting findings: that workers seek to avoid responding to social incentives when financial incentives are strong; that there is heterogeneity in how susceptible individuals are to social effects and that those who are able to avoid them are more likely to advance to elite professional labor markets; and that workers learn with professional experience not to be affected by social forces. We view our results as complementary to the existing studies of peer effects in the workplace and as a first step towards explaining how these social effects vary across labor markets, across individuals and with changes in the form of incentives faced. In addition to the empirical results on peer effects in the workplace, we also point out that many typical peer effects regressions are biased because individuals cannot be their own peers, and suggest a simple correction.

Jonathan Guryan
University of Chicago
Graduate School of Business
5807 S. Woodlawn Ave.
Chicago, IL 60637
and NBER
jonathan.guryan@chicagoGSB.edu

Matt Notowidigdo
MIT Department of Economics
E52-204F
50 Memorial Drive
Cambridge MA 02142
noto@mit.edu

Kory Kroft
University of California, Berkeley
Department of Economics
549 Evans Hall #3880
Berkeley CA 94720-3882
kroft@econ.berkeley.edu

1 Introduction

Is an employee's productivity influenced by the productivity of his or her nearby co-workers? The answer to this question is important for the optimal organization of labor in a workplace and for the optimal design of incentives.^{1 2} Despite the importance of this question, empirically it is difficult to address. The main difficulty is that it is hard to disentangle whether an observed correlation in behavior is due to the effect of the group's behavior on individual behavior ('endogenous effects'), the effect of the group's characteristics on individual's behavior ('contextual effects'), or the correlation between observed and unobserved determinants of the outcome ('correlated effects').³ Only the first two of these effects represent true peer effects.⁴

Despite this identification problem, a few empirical studies have found evidence that peer effects are important in several workplace settings: among low-wage workers at a grocery store (Mas and Moretti, 2006), among soft-fruit pickers (Bandiera, Barankay and Rasul, 2007), and among workers performing a simple task in a laboratory setting (Falk and Ichino, 2006). The present study contributes to this literature by estimating the importance of peer effects among elite professionals: male professional golfers on the PGA Tour. Moreover, this is the first field study of peer effects in the workplace that has explicit random assignment of peers.

In the first two rounds of PGA tournaments, players are randomly assigned to groups of three conditional on their *category*, which takes one of three values and is determined using a simple formula that depends on past performance. Within a playing group, players are proximate to one another and can therefore observe each others' shots and scores. This proximity creates the

¹The relevance for the optimal design of incentives hinges on whether social effects are complements or substitutes for financial incentives.

²Complementarities between an employee's productivity and the productivity of his peers may arise for at least three reasons: (1) individuals may learn from their co-workers about how best to perform a given task, (2) workers may be motivated to exert effort when they see their co-workers working hard or performing well or when they know their co-workers are watching, or (3) the nature of the production process may be such that the productivity of one worker mechanically influences the productivity of another worker directly (e.g. the assembly line). While the former two sources are 'behavioral', the latter is a mechanical effect that arises for purely structural or technological reasons. For clarity, we label the first two of these effects 'peer effects' and the last effect a 'production complementarity'.

³For example, members of a group might be hit by a common shock that affects behavior independent of spillovers or social interactions.

⁴See Manski (1993) and Moffitt (2001) for descriptions of the problems associated with estimating peer effects. An additional difficulty is what Manski calls the 'reflection problem'. To understand the reflection problem, consider trying to estimate the effect of individual i 's behavior on individual j 's behavior. It is very difficult to tell which member of the pair is affecting the other's behavior, and whether the affected behavior by one member of the pair affects the other's in turn, and so on.

opportunity to learn from and be motivated by peers.

There are several learning opportunities during a round. For example, a player must judge the direction of the wind when hitting his approach shot to the putting green. Wind introduces uncertainty into shot and club selection. Thus, by observing the ball flight of others in the playing group, a player can reduce this uncertainty and increase his chance of hitting a successful shot. Another example is the putting green. Subtle slopes, moisture, and the type of grass all affect the direction and speed of a putt. The chance to learn how a skilled putter manages these conditions may confer an advantage to a peer golfer.

Turning to motivation, there are several ways that motivation can affect performance. The chance to visualize a good shot may help a player to execute his own shot successfully. Similarly, seeing a competitor play well may directly motivate a player and help him to focus his mental attention on the task at hand. Still more, a player's self-confidence, and in turn his performance (Benabou and Tirole, 2002), may be directly affected by the abilities or play of his peers (Festinger, 1954).

In the parlance of the social interactions literature, this playing group is a player's reference group and within it a player's actions may be influenced by the group's actions and/or characteristics. We construct measures of player ability using information on past performance and we test for peer effects by considering whether (random) variation in peers' ability affects own performance.

The main result of this paper is that neither the ability nor the current performance of playing partners affect the performance of professional golfers. In our preferred specification, we can rule out peer effects larger than 0.045 strokes for a 1 stroke increase in playing partners' ability and our point estimate is actually negative. The results are robust to alternative peer effects specifications. We offer several explanations for our contrasting findings: (1) that workers seek to avoid the social effects found in previous studies when there are strong performance-based financial incentives; (2) that there is heterogeneity in how susceptible individuals are to social effects and that those who are able to avoid them are more likely to advance to elite professional labor markets; and (3) that workers learn with professional experience not to be affected by social forces. We view these results as complementary to the existing studies of peer effects in the workplace, a first step towards explaining how these social effects vary across labor markets, across individuals and with changes in the form of incentives. That there might be heterogeneity in peer effects is

particularly important in light of the fact that previous studies typically focus on low-skill jobs or take place in the laboratory. If we are correct that high-skilled professionals are less susceptible to social pressures at work, this finding is relevant for large parts of the labor market for which there currently exists no evidence on peer effects (e.g. financial professionals, lawyers, doctors, professors).

Our study contributes to the literature on peer effects in the workplace in six important ways. First, we make a methodological point concerning a bias inherent in most typical peer effects regressions and suggest a simple correction. Because individuals cannot be their own peers, even random assignment generates a negative correlation in pre-determined characteristics of peers. Intuitively, the urn from which the peers of an individual are drawn does not include the individual. Thus, the population at risk to be peers with high-ability individuals is on average lower ability than the population at risk to be peers with low-ability individuals. As a result, the typical test for random assignment, a regression of i 's predetermined characteristic on the mean characteristic of i 's peers, produces a slightly negative coefficient even when peers are truly randomly assigned. Because this mechanism causes negative correlation in both observed and unobserved characteristics of peers, most regression-based tests of peer effects are affected by it. We present results from Monte Carlo simulations showing that this bias can be reasonably large, and that it is decreasing in the size of the population from which peers are selected. We then propose a simple solution to this problem—controlling for the average ability of the population at risk to be the individual's peers—and show that including this additional regressor produces test statistics that are well-behaved.

Second, our research design—the random assignment of peers—allows us to estimate the causal effect of a playing partners' ability on own performance. To our knowledge, ours is the first field study of peer effects in the workplace to have explicit random assignment of co-workers.

Third, the design of golf tournaments allows us to test directly for the most likely sources of common shocks. As pointed out by Manski (1993), even with random assignment a correlation of own outcome with the outcome of the group is not informative on the importance of endogenous peer effects if there are common unobserved shocks. We use the performance of nearby groups, those who are playing a few holes ahead or behind the reference group, to measure the role of common shocks (e.g. from common weather conditions which vary over the course of the day).

Fourth, our results are purged of ‘relative performance effects’ because the objective for each player in a golf tournament is to score the lowest, regardless of with whom that player is playing.⁵ Pay is based on relative performance, but performance is compared to the entire field of entrants in a tournament, not relative to the players within a playing group. This contrasts with other settings, such as classrooms, where the performance of an individual is assessed relative to the individual’s peers. For example, it tells us nothing about whether students learn from smarter peers if individuals try harder because of incentives created by relative performance evaluation in the classroom. In a golf tournament, there is no reason (other than better information about how the player next to him is doing) for a player to care more about the performance of his playing partner than about anyone else in the tournament.⁶

Fifth, the set-up of golf tournaments allows us to identify peer effects in a setting devoid of most production-technology complementarities. This is unlike, for example, the grocery store setting considered in Mas and Moretti (2006), where a grocery store scanner’s productivity may depend on the productivity of nearby scanners for the simple reason that a customer’s incentive is to choose the fastest aisle.

Sixth, though we present a basic result based on an overall measure of skill, we have multi-dimensional measures of ability that line up nicely with important potential underlying mechanisms of peer effects. In particular, we are in principle able to distinguish between learning effects and motivational effects, the latter of which are more psychological in nature.

The remainder of this paper is organized as follows: section 2 reviews the literature, section 3 briefly describes the PGA Tour, section 4 describes the methodological point concerning bias in typical peer effects regressions, section 5 shows results verifying the random assignment mechanism, section 6 discusses our empirical approach, section 7 discusses the validity of our empirical strategy and our results and section 8 concludes.

⁵This objective is accurate for almost every player, except perhaps the players who have a reasonable probability of earning the top few prizes. In those cases the objective is actually to score lower than your opponents (where your opponents are the small universe of players who are competing with you for the top prizes). With this concern in mind, in our results section we drop top-tier players as a robustness check which does not change our main results.

⁶The scores of the current tournament leaders are typically posted around the course so that golfers have information about how they are playing relative to players outside of their group. This information makes the information about how their playing partners are performing significantly less valuable.

2 Related Literature

It has long been recognized by psychologists that an individual's performance might be influenced by his peers. The first study to show evidence of such peer effects was Triplett (1898), who noted that cyclists raced faster when they were pitted against one another, and slower when they raced only against a clock. While Triplett's study shows that the presence of others can *facilitate* performance, others found that the presence of others can *inhibit* performance. In particular, Allport (1924) found that people in a group setting wrote more refutations of a logical argument, but that the quality of the work was lower than when they worked alone. Similarly Pessin (1933) found that the presence of a spectator reduced individual performance on a memory task. Zajonc (1965) resolved these paradoxical findings by pointing out that the task in these experimental setups varied in a way that confounded the results. In particular, he argued that for well-learned or innate tasks, the presence of others improves performance. For complex tasks however, he argued that the presence of others worsens performance.

Guided by the intuition that peers may affect behavior and hence market outcomes, several economic studies of peer effects have recently emerged in a variety of domains. Examples include education (Graham, 2004), crime (Glaeser, Scheinkman and Sacerdote, 1996), unemployment insurance take-up (Kroft, 2007), welfare participation (Bertrand, Luttmer and Mullainathan, 2000), and retirement planning (Duflo and Saez, 2004). The remainder of this section reviews three papers that are conceptually most similar to our research, in that they attempt to measure peer effects in the workplace or in a work-like task.

The first economic study of peer effects in a work-like setting is a laboratory experiment conducted by Falk and Ichino (2006). This experiment measures how an individual's productivity is influenced by the presence of another individual working on the same task: stuffing letters into envelopes. They find moderate and significant peer effects: a 10% increase in peers' output increases a given individual's effort by 1.4%. A criticism of this study is one that applies broadly to other studies in the lab; in particular, that it may have low external validity because of experimenter demand effects or because experimental subjects get paid minimal fees to participate and as a result their incentives may be weak (Levitt and List, 2007).⁷

⁷Another concern about laboratory experiments is that subjects are prevented from sorting into environments based on their social preferences (Lazear, Malmendier, and Weber, 2006).

Two recent studies of peer effects in the workplace have examined data collected from the field. Mas and Moretti (2006) measure peer effects directly in the workplace using grocery scanner data. There is not explicit randomization, but the authors present evidence that the assignment of workers to shifts appears haphazard. Since grocery store managers do not measure individual output directly in a team production setting, one might expect to observe significant free riding and sub-optimal effort. Instead, this study finds evidence of significant peer effects with magnitudes similar to those found by Falk and Ichino (2006): a 10% increase in the average permanent productivity of co-workers increases a given worker's effort by 1.7%. As Mas and Moretti discuss in their paper, grocery scanner is an occupation where compensation is not very responsive to changes in individual effort and output. They conjecture that "economic incentives alone may not be enough to explain what motivates [a] worker to exert effort in these jobs."

Bandiera, Barankay and Rasul (2007) examine how the identity and skills of nearby workers affect the productivity of soft-fruit pickers on a farm. Assignment of workers to rows of fruit is made by a combination of managers. Though assignment is not explicitly random, the authors present evidence to support the claim that it is orthogonal to worker productivity. The authors find that productivity responds to the presence of a friend working nearby, but do not show evidence that productivity responds to the skill-level of non-friend co-workers. High-skilled workers slow down when working next to a less productive friend, and low-skilled workers speed up when working next to a more productive friend. Workers are paid piece rates in this setting and are willing to forgo income to conform to a social norm along with friends. The authors also find that workers respond to the cost of conforming: on high-yield days the relatively high-skilled friend slows down less and the relatively low-skilled friend works harder.

In light of the fact that monetary incentives are weaker in the Mas and Moretti (2006) study than in Bandiera, et al.'s (2007), it is interesting to note that Mas and Moretti (2006) find more general peer effects. A recent paper by Lemieux et al. (2006) points out that an increasing fraction of US jobs contain some type of performance pay based either on a commission, a bonus, or a piece rate.⁸ Since neither of the aforementioned studies contain very strong economic incentives, it is natural to ask whether peer effects also exist in settings that do.⁹ In the setting that we consider—

⁸Using the CPS, the authors find that the overall incidence of performance pay was a little more than 30 percent in the late 1970's but grew to over 40 percent by the late 1990's.

⁹In a prior paper, Bandiera, Barankay and Rasul (2005) find an increase in effort in response to a switch in com-

professional golf tournaments—compensation is determined purely by performance.¹⁰ Pay is high and the pay structure is quite convex. For example, during the 2006 PGA season, the top prize money earner, Tiger Woods, earned almost \$10 million, and 93 golfers earned in excess of \$1 million during the season. We discuss the convexity of tournament payouts later in the paper.

The existing field studies in the literature on peer effects in the workplace have focused on low-skilled jobs in particular industries. It is possible that there is heterogeneity in how susceptible individuals are to social effects at work. Motivated by this, we ask whether peer effects exist in workplaces made up of highly skilled professional workers.¹¹

In terms of research design, our study is related to Sacerdote (2001) and Zimmerman (2003), who measure peer effects in higher education using the random assignment of dorm roommates at Dartmouth and Williams Colleges, respectively. Sacerdote finds that an increase in roommate's GPA by 1 point increases own freshman year GPA by 0.120. The coefficient on roommate's GPA, however, drops significantly—from 0.120 to 0.068—when dorm fixed effects are included, suggesting that common shocks might be driving some of the correlation in GPAs between roommates. As we discuss below, we are able to test directly for the most likely source of common shocks to golfers.¹²

3 Background and Data

3.1 Institutional Details of the PGA Tour

*Relevant Rules.*¹³ There is a prescribed set of rules to determine the order of play. The player with the best (i.e. lowest) score from the previous hole takes his initial shot first, followed by the player with the next best score from the previous hole. After that, the player who is farthest from the hole

pensation regime from relative pay to piece rate pay. Their evidence suggests that social preferences can be offset by appropriate monetary incentives.

¹⁰Golfers also receive a sizeable amount from professional endorsements. Presumably, endorsement earnings are related indirectly to performance. Tiger Woods, the golfer with the highest earnings from tournaments is also the golfer with the greatest endorsement income.

¹¹In Section 7.5, we will also look at whether there are heterogeneous peer effects within this class of professional workers.

¹²Finally, other studies have used professional golf data to test economic theories. Bognanno and Ehrenberg (1990) test whether professional golf tournaments elicit effort responses. They find that the level and structure of prizes in PGA tournaments influence players' performance. However, Orszag (1994) shows their results might not be robust once weather shocks are accounted for.

¹³Appendix A gives a short introduction to the basic rules of golf.

always shoots next.¹⁴

Selection. Golfers from all over the world participate in PGA tournaments that are held (almost) every week.¹⁵ At the end of each season, the top 125 earners in PGA Tour events are made full-time members of the PGA Tour for the following year. Those not in the top 125 and anyone else who wants to become a full-time member of the PGA Tour must go to ‘Qualifying School’, where there are a limited number of spots available to the top finishers. For most PGA Tour tournaments, the players have the right but not the obligation to participate in the tournament. In practice, there is variation in the fraction of tournaments played by PGA Tour players. The median player plays in about 59 percent of a season’s tournaments.¹⁶ Some players avoid tournaments that do not have a large enough purse or a high prestige, while other players might avoid tournaments that require a substantial amount of travel. Bronars and Oettinger (2001) look directly at several determinants of selection into golf tournaments; they find a substantial effect of the level of the purse on entry decisions. Most relevant for our purposes is the fact that because membership on the tour for the following year is based on earnings, lower-earning players have an incentive to enter tournaments that higher-skilled players choose not to enter. It is therefore necessary to condition on tournament-by-category fixed effects in the empirical work since random assignment takes place at this level. Conditional on the set of players who enter a tournament, playing partners are randomly assigned within categories. Unconditional on this fully interacted set of fixed effects, assignment is not random.

Tournament Details. Tournaments are generally four rounds of 18 holes played over four days, and prizes are awarded based on the cumulative performance of the players over all four rounds.

¹⁴In general we expect better players to be more likely to shoot first on the initial shot of a hole, but worse players are more likely to shoot first on subsequent shots on the hole. This offsetting pattern leads us to believe that the order of play does not significantly affect the peer effects we estimate. If one could collect shot-by-shot data, it would be interesting to examine whether behavior is affected by players that shoot immediately prior.

¹⁵In 2007, there were 47 PGA tournaments played over 44 weeks. The three instances where there are two tournaments during a week are the ‘major championships.’ Because these championships are only for qualifying golfers from around the world (and not exclusively for PGA Tour members), and because these tournaments are not sponsored by the PGA Tour, the PGA Tour also hosts tournaments during the major championships for the remaining PGA Tour members who did not qualify for the major tournaments. We drop all the major championships from our data set because they do not use the same random assignment mechanism.

¹⁶Since most players play in most tournaments, when we construct our measure of ability we will usually have enough past tournament results to reliably estimate the ability of each player. The median player has 30 past tournament results with which to estimate his ability, 29 percent of the players have more than 40 past tournaments, and 14 percent of the players have fewer than 5 tournaments. We weight all regressions by the number of past performance observations used to compute ability.

At the end of the second round, there is a ‘cut’ that eliminates approximately half of the tournament field based on cumulative performance over the first two rounds. Most tournaments have 130-160 players, and the cut reduces the field to the top 70 players (plus ties) after the first two rounds.

Economic Incentives. In order to earn prize money, players must survive the cut.¹⁷ The prize structure is extremely convex: first prize is generally 18 percent of the total purse. Also, because better performance will also likely lead to endorsement money, the true economic incentives are even stronger. Figure 1 shows the convexity in the prize structure of a typical tournament, and Table 2 shows the distribution of total purses in our sample. Below we examine whether the tournament prize structure and purse size affect player performance.

Assignment Rule. Players are continuously assigned to one of three *categories* according to rules described in detail in Appendix B. Players in category 1 are typically tournament winners from the current or previous year or players in the top 25 on the earnings list from the previous year. These include players such as Tiger Woods, Phil Mickelson and Ernie Els. Players in category 2 are typically those between 26 and 125 on the earnings list from the previous year, and players who have made at least 50 cuts during their career or who are currently ranked among the top 50 on the World Golf Rankings. This group is the largest in most tournaments and includes players such as Nick Faldo and John Daly. Category 3 consists of all other entrants in the tournament. Within these categories, tournaments then randomly assign playing partners to groups of three golfers.¹⁸ These groups then play together for the first two rounds of the tournament. We evaluate performance from the first two rounds only, since in the third and fourth rounds players are assigned based on performance of the previous rounds.

3.2 Data

Key Variables. We collected information on tee times, groupings, results, earnings, course slope and rating, and player statistics and characteristics from the PGA Tour website and various other

¹⁷There is no entry fee for PGA Tour members to play in PGA tournaments. Non-members must pay a nominal \$400 entry fee.

¹⁸This is similar to the random assignment mechanism used to assign roommates at Dartmouth, as discussed in Sacerdote (2000)—conditional on an observable characteristic, there is random assignment within individuals sharing that characteristic.

websites.¹⁹ Most of our data spans the 1999-2006 golf seasons; however, we only have tee times, groupings (i.e. the peer groups) and categories for the 2002, 2005, and 2006 seasons.²⁰ As discussed below, to construct a pre-determined ability measure, we use the 1999, 2000, and 2001 data for the players from the 2002 season and the 2003 and 2004 for the players from the 2005 and 2006 seasons.²¹ Our data therefore allow us to observe the performance of the same individuals playing in many PGA tournaments over three seasons (2002, 2005, and 2006).

We also collected a rich set of disaggregated player statistics for all years of the sample. These variables were collected in hopes of shedding light on the mechanism through which the peer effects operate. These measures of skill include the average number of putts per round, average driving distance, and a measure of accuracy, the average number of greens hit in regulation (the fraction of times a golfer gets the ball onto the green in at least two shots less than par). We discuss below how these measures might allow us to separately identify two specific forms of peer effects: learning and motivation.

Sample Selection. As discussed in more detail in Appendix B, the scores from some tournaments must be dropped because random assignment rules are not followed. The four most prestigious tournaments (called ‘major championships’) do not use the same conditionally random assignment mechanism. For example, the U.S. Open openly admits to creating ‘compelling’ groups to stimulate television ratings. The vast majority of PGA tournaments, however, use the same random assignment mechanism, and we have confirmed this through several personal communications with the PGA Tour, and through statistical tests reported below.

3.3 A Measure of Ability

In order to estimate peer effects, we require a measure of ability or skill for every player. We construct this measure by averaging playing scores in prior years.²² There is a problem with using a player’s raw scoring average as a proxy for ability, however. There is heterogeneity in golf

¹⁹We gathered data from www.pgatour.com, www.espn.com, www.usatoday.com, and www.cnnsi.com, and www.yahoo.com.

²⁰Tee times were collected each Thursday during the 2002 season since a historical list was not maintained either on web sites or by the PGA Tour at that time. The 2005 and 2006 tee times were collected subsequently in an effort to increase sample size and power.

²¹We use only the first two rounds to construct our measure of pre-determined ability.

²²Since our peer effects specifications only use the first two rounds of a tournament, we construct our ability measure using only the first two rounds from earlier tournaments.

course difficulty and better players tend to self-select into tournaments with harder courses. As a result, differences in measured scoring average tend to understate differences in true ability.

We address this problem using a simplified form of the official handicap correction used by the United States Golf Association (USGA), the major golf authority that oversees the official rules of the game. The USGA measures the difficulty for most golf courses in the United States. Using data on scores from golfers of different skill levels, the USGA assigns a *slope* and a *rating* which can respectively be thought of as related to the estimated slope and intercept from a regression of score on ability.²³ We adjust the scores using the slope and rating and take the average of these adjusted scores from prior years for each player.²⁴ We have also experimented with several other measures of ability including using the (raw) past scoring average and a best linear predictor, and we have found qualitatively similar results.

4 Bias in Typical Peer Effects Tests

4.1 Explaining the problem

Before presenting the empirical results, in this section we describe an important methodological consideration. Given the importance of random assignment, papers that report estimates of peer effects typically present statistical evidence to buttress the case that assignment of peers is random, or as good as random. The typical test is an OLS regression of individual i 's predetermined characteristic x on the average \bar{x} of i 's peers, conditional on any variables on which randomization was conditioned. The argument is made that if assignment of peers is random, or if selection into peer groups is ignorable, then this regression should yield a coefficient of zero. In our case, this regression would be of the form

$$Ability_{ikt} = \pi_1 + \pi_2 \times \overline{Ability}_{-i,kt} + \delta_{tc} + \epsilon_{ikt} \quad (1)$$

²³To compute a golfer's handicap, the USGA first computes the handicap differential for each of the golfer's last 20 scores according to the formula: handicap differential = (score-rating)*(113/slope). The USGA normalizes slope such that 113 is the slope of a course of 'standard difficulty'. The handicap index is then calculated as 0.96 times the average of the lowest 10 handicap differentials from the most recent 20 scores.

²⁴This differs from the official handicap measure used by the USGA in one important way. The official measure is based on the best ten of the players most recent twenty completed rounds. The USGA measure is intended to be used to allow amateur golfers of differing abilities to compete against each other. Since we are interested in a measure of skill, we average over all prior rounds in our sample.

where i indexes players, k indexes (peer) groups, t indexes tournaments, c indexes categories, and δ_{tc} is a fully interacted set of tournament-by-category dummies, including main effects. This is, for example, the test for random assignment reported in Sacerdote (2001).

This test for the random assignment of individuals to groups is not generally well-behaved. The problem stems from the fact that an individual cannot be assigned to himself. In a sense, sampling of peers is done without replacement—the individual himself is removed from the ‘urn’ from which his peers are chosen. As a result, the peers for high-ability individuals are chosen from a group with a slightly lower mean ability than the peers for low-ability individuals.

Consider an example in which four individuals are randomly assigned to groups of two. To make the example concrete, let the individuals have pre-determined abilities 1, 2, 3, and 4. If pairs are randomly selected, there are three possible sets of pairs. Individual 1 has an equal chance of being paired with either 2, 3, or 4. So, the *ex-ante* average ability of his partner is 3. Individual 4 has an equal chance of being paired with either 1, 2, or 3, and thus the *ex-ante* average ability of his partner is 2. This mechanical relationship between own ability and the mean ability of randomly-assigned peers—which is a general problem in all peer effects studies—causes estimates of equation (1) to produce negative values of $\hat{\pi}_2$. Random assignment appears non-random, and positively matched peers can appear randomly matched.

This bias is decreasing in the size of the population from which peers are drawn, i.e. the size of the ‘urn’. As the urn increases in size, each individual contributes less to the average ability of the population from which peers are drawn, and the difference in average ability of potential peers for low and high ability individuals converges to zero. In settings where peers are drawn from large groups, ignoring this mechanical relationship is inconsequential. In our case, the average urn size is 60, and 25 percent of the time the urn size is less than 18.

We present Monte Carlo results in Figure 2 which confirm that estimates of (1) are negatively biased and that the bias is decreasing in the size of the urn. We report the results from two simulations. For the first simulation, we created 55 players with ability drawn from a normal distribution with mean zero and standard deviation one. We then created 100 tournaments and for each tournament randomly selected M players. Each of these M players were then assigned to groups of three. M , which corresponds to the size of the urn from which peers were drawn, was randomly chosen to be either 39, 42, 45, 48, or 51 with equal probability for an average M of 45.

We explain below the reason for the variation in M . Finally, we estimated an OLS regression of own ability on the average of partners' ability, controlling for tournament fixed effects, and the estimate of $\hat{\pi}_2$ and p -values were saved. This procedure was repeated 10,000 times. For the second simulation, we increased the average size of M by creating 550 players, and allowing M to take on values of 444, 447, 450, 453, and 456 with equal probability for an average M of 450.

The results from the first simulation are shown on the left side of Figure 2. As predicted, the typical OLS randomization test is not well behaved. The test substantially overrejects, rejecting at the 5-percent level more than 18 percent of the time. Even though peers are randomly assigned, the estimated correlation between abilities of peers is on average -0.046. This negative relationship is exactly as one should expect, resulting from the fact that individuals cannot be their own peers.

The intuitive argument made above also implies that the size of the bias should be decreasing in M . Indeed, this is the case. The right side of Figure 2 shows results from the second simulation where the urn size was increased by an order of magnitude. The typical randomization test is more well-behaved. The estimates of $\hat{\pi}_2$ center around zero, the test rejects at the 5-percent level 5.4 percent of the time, and p -values are close to uniformly distributed between 0 and 1. In short, the Monte Carlo results show that the typical test for randomization is biased when the set of individuals from which peers are drawn is relatively small.

Bias stemming from this problem is not limited to tests of random assignment. The argument above applies to unobserved characteristics just as it does to observed characteristics. Because no one can be his own peer, individuals with high unobservables are at risk to be peers with individuals who have lower average unobservables than individuals with low unobservables are. Thus, regressions of individual i 's *outcome* on the average characteristics or outcomes of his peers are also negatively biased. F -tests on sets of indicator variables for being peers with individual j are also biased for the same reason. Being paired with individual j is informative about individual i 's observables and unobservables.

4.2 A proposed solution

To our knowledge, this point has not been made clearly in the literature.²⁵ We propose a simple correction to equation (1) that will produce a well-behaved test of random assignment of peers, even with small urn sizes. Since the bias stems from the fact that each individual's peers are drawn from a population with a different mean ability, we simply control for that mean. Specifically, we add to equation (1) the mean ability of all individuals in the urn, excluding individual i . The modified estimating equation is thus

$$Ability_{ikt} = \pi_1 + \pi_2 \times \overline{Ability}_{-i,kt} + \delta_{tc} + \phi \times \overline{Ability}_{-i,ct} + u_{ikt} \quad (2)$$

where $\overline{Ability}_{-i,ct}$ is the mean ability of all players in the same category \times tournament cell as player i , other than player i himself (i.e. all individuals that are eligible to be matched with individual i), and ϕ is a parameter to be estimated. It should be noted that it is necessary for there to be variation in the set of players in player i 's urn to be able to estimate this regression.²⁶

Figure 3 shows the results from Monte Carlo simulations analogous to those reported above. The difference here is that instead of estimating the typical OLS regression, we include $\overline{Ability}_{-i,ct}$ as an additional regressor. As can be seen clearly in the figure, the addition of this control makes the OLS test of randomization well-behaved regardless of whether average urn size is large or small. In both cases, the estimated correlation of peers ability centers around zero, the test rejects at the 5-percent level approximately 5 percent of the time, and p -values are approximately uniformly distributed between 0 and 1. In results not reported here we have also confirmed that the test can detect deviations from random assignment. Going forward, we therefore include $\overline{Ability}_{-i,ct}$ as a control in all specifications, both tests for random assignment and tests of peer effects.

²⁵The closest discussion of which we are aware is by Boozer and Cacciola (2001), who point out that the ability to detect peer effects in a linear-in-means regression of outcomes on mean outcomes is related to the size of the reference group, but they do not link this discussion to tests for random assignment nor to the fact that individuals cannot be assigned to themselves as peers.

²⁶If every urn has N players, then my ability $Ability_{iktr}$ is related to the mean ability in my urn $\overline{Ability}_{ct}$ and the 'leave-me-out' mean $\overline{Ability}_{-i,ct}$ by the following identity: $Ability_{iktr} = N * \overline{Ability}_{ct} - (N - 1) * \overline{Ability}_{-i,ct}$.

5 Verifying Random Assignment of Peers

Using our various measures of ability, we now test the claim that assignment to playing groups is random within a tournament conditional on the player's category.²⁷ In this section, we report results from estimating variations of specification (2), with various measures of pre-determined ability. Before reporting those results, however, we first report the results of estimating equation (2) with only tournament fixed effects (i.e. dropping category fixed effects and tournament-by-category interactions). We do this to show that the randomization test has sufficient power to detect deviations from random assignment in a setting where we know assignment is not random. The results are shown in column (1) of Table 1. The coefficient on average partners' ability is 0.159 and is strongly statistically significant. In column (2) we show results from estimating equation (2) with only category fixed effects (i.e dropping tournament fixed effects and tournament-by-category interactions): the coefficient on average partners' ability is 0.116 and is strongly statistically significant. The former result is consistent with the fact that players are not *unconditionally* randomly assigned within a tournament. Stratification by category means that better players are grouped with better players. The latter indicates that even conditional on category, players of similar abilities select into similar tournaments, which is exactly what we would expect if lower-ability players strategically choose tournaments that higher-ability players avoid, since pay within tournaments is based on performance relative to the tournament's actual entrants (Bronars and Oettinger, 2001). Importantly, both results show that the test is powerful enough to reject random assignment in a situation where we know assignment was not conditionally random.

We next present the results of the correct randomization test, which includes the full set of tournament-by-category fixed effects along with the control for $\overline{Ability}_{-i,ct}$. These results are shown in column (3). The estimate of the correlation between own ability and playing partners' ability is small and insignificant. This is just as is implied by the random assignment mechanism that the PGA Tour claims to use. We also test for random assignment using various disaggregated measures of ability (e.g. driving distance, putts per round and greens in regulation per round). If players are conditionally randomly assigned to groups, these measures should not be correlated conditional on the correct set of variables described above. The results of estimating equation

²⁷This claim is based on the PGA Player Handbook and Tournament Regulations, and on numerous telephone conversations with PGA Tour officials.

(1), replacing average adjusted score with these measures, are shown in columns (4), (5), and (6) of Table 1. The coefficients on the other characteristics are also statistically insignificant and the point estimates are small. Lastly, columns (7)-(9) separately estimate the randomization equation by category, which also shows no evidence of deviations from conditional random assignment. Overall, we interpret the results of Table 1 as supporting the claim that the PGA Tour assigns players randomly to groups conditional on categories.

6 Empirical Framework

To estimate peer effects, the data are analyzed using a simple linear model where own score depends on own ability and playing partners' ability. The key identifying assumption is that, conditional on tournament and category, players are randomly assigned to groups.²⁸ Thus by controlling for tournament-specific category fixed effects, we can estimate the causal effect of playing partners' ability on own score. Our baseline specification is

$$Score_{iktr} = \alpha_1 + \beta_1 \times Ability_i + \gamma_1 \times \overline{Ability}_{-i,kt} + \delta_{tc} + \varphi_1 \times \overline{Ability}_{-i,ct} + e_{iktr} \quad (3)$$

where i indexes players, k indexes groups, t indexes tournaments²⁹, r indexes each of the first two rounds of a tournament, c indexes categories, δ_{tc} is a full set of tournament-by-category dummies to be estimated, α_1 , β_1 , γ_1 , and φ_1 are parameters, and e is an error term. The parameter γ_1 measures the effect of the average ability of playing partners on own score, and is our primary measure of peer effects. Its magnitude is generally evaluated in relation to the magnitude of β_1 , which is the effect of own ability on own score. Since playing partners are randomly assigned, the coefficients in equation (3) can be estimated consistently using OLS. The parameter φ_1 is the same 'leave-me-out' mean correction described in the previous section. Because we included it in the randomization test, we also include it here, although it does not affect any of our results.

Even with the handicap correction described above, a remaining potential problem with our measure of ability is measurement error. This should be a concern for all studies of peer effects

²⁸This is the assumption that was tested empirically in section 5.

²⁹For the remainder of this paper, we define a 'tournament' to be a tournament-by-year cell, since our dataset has several tournaments that are played again in subsequent years. For example, the Ford Championship in 2002 has a separate dummy than the Ford Championship in 2005.

in the workplace that estimate exogenous effects. A nice feature of the specification above is that it contains a simple correction for measurement error. Because they are measured using the same variable, there should be the same degree of measurement error in each golfer's individual measure of ability. If each golfer had a single peer, the estimates of $\hat{\beta}_1$ and $\hat{\gamma}_1$ would be equally attenuated. The ratio $\hat{\gamma}_1/\hat{\beta}_1$ would in that case give us a measurement-error-corrected estimate of the reduced-form exogenous peer effect. Because in our case $\overline{Ability}_{-i,kt}$ is an average of two values, it is likely to be less error-ridden. We therefore take $\hat{\gamma}_1/\hat{\beta}_1$ to be an upper bound of the measurement-error-corrected estimate of the reduced-form exogenous peer effect.

A key advantage to estimating the reduced-form specification in (3) is that the average ability of playing partners is a pre-determined characteristic. Thus, our estimate of γ_1 is unlikely to be biased due to the presence of common unobserved shocks.³⁰ An alternative commonly estimated specification replaces peers' ability with peers' score. This outcome-on-outcome specification estimates a combination of endogenous and contextual effects, but intuitively examines how performance relates to the contemporaneous performance of peers, rather than just to peers' predetermined skills. Even with random assignment, however, one cannot rule out that a positive relationship between own score and peers' score is driven by shocks commonly experienced by individuals within a peer group. We nevertheless run regressions of the following form to get an upper bound on the magnitude of peer effects:

$$Score_{iktr} = \alpha_2 + \beta_2 \times Ability_i + \gamma_2 \times \overline{Score}_{-i,kt} + \delta_{tc} + \varphi_2 \times \overline{Ability}_{-i,ct} + v_{iktr} \quad (4)$$

where $\overline{Score}_{-i,kt}$ is the average score in the current round of player i 's playing partners. Because common shocks are expected to cause upward bias, the estimate of γ_2 should be viewed as an upper bound on the extent of peer effects. A nice feature of our setup also allows us to gauge the magnitude of the bias created by common shocks since we can observe the scores of nearby playing groups, who are likely to be affected by similar shocks. We report estimates that take advantage of this feature of the research design in the following section.

³⁰Note however that these common shocks are likely to affect the standard errors. Hence in the peer effect regressions below, we cluster at the group level. The estimated standard errors are virtually unchanged if we cluster by tournament \times category, and are actually smaller if we cluster by tournament.

7 Results

7.1 Summary Statistics

Before turning to the main results of the paper we first present the descriptive statistics in Table 2. It is important to understand the units of our primary variables of interest. Score is a variable that represents the actual golf score the player achieved in a given tournament-round. Ability is in the same units as Score (i.e. golf strokes). Usually the player's score is measured relative to the par on the course, which is typically 72 strokes. Ability, while in the same units as Score, is typically expressed as deviation of score from par. Throughout the results section, it is helpful to keep in mind that lower scores in golf indicate better performance (and, analogously, a lower Ability measure indicates a higher ability player).

Figures 4 and 5 show the distribution of handicap both in the full sample, and broken down by categories. Two things should be noted from these figures. First, there is a reasonable amount of variation in ability even among professional golfers. The difference between the 90th percentile and 10th percentile in adjusted average score—our baseline measure of ability—is 1.65. Perhaps more importantly, given the stratification by category prior to random assignment, is that much of the variance in ability remains after separating by categories. As can be seen clearly in Figure 5, the average scores increase monotonically from category 1 to category 3. However, there is a great deal of overlap in the distributions. The 90-10 differences in measured ability in category 1, 2, and 3 are 1.71, 1.38, and 2.38 strokes respectively. These differences represent wide ranges in ability — using our data on earnings, a reduction in handicap of 1.65 translates into an increase in expected earnings of 70 percent.³¹ This suggests that differences in strokes on this order of magnitude should be quite salient to players.

There is also a good deal of variation across players in specific dimensions of past performance. The 90th percentile golfer drives the ball 24 yards farther than the golfer at the 10th percentile in average driving distance. This is 8.6 percent of the mean drive. For putting, the differences are similar: the 90th percentile putter hits 2.4 fewer putts per round than the 10th percentile putter (8.3 percent of the mean). And for accuracy, the 90-10 difference is 12.6 percent of the mean for

³¹A regression of log earnings on handicap with a full set of tournament fixed effects, category fixed effects and their interactions gives a coefficient on handicap of -0.320 with a t-statistic of 13.24. Thus a reduction in handicap of 1.65 will increase earnings by 0.528 log points, or 69.6%.

greens hit in regulation.

7.2 Visual Evidence

To get a sense of the importance of peer effects, we first plot adjusted scores against playing partners' handicap. To do this, we first regress each player's score on tournament and category fixed effects and their interactions and the 'leave-me-out' mean of the tournament-by-category urn. Then we take the residuals from this regression, compute means by each decile of the partners' ability distribution, and graph the average residual against each decile bin. Figure 6 reports this graph for the full sample, which shows zero correlation between own score and the ability of randomly assigned playing partners. Those who were randomly assigned to partners with higher average scores scored no differently than those who were assigned to partners with low average scores. There also does not appear to be evidence of non-linear peer effects. We take this to be a first piece of evidence that peer effects among professional golfers are economically insignificant. To place a confidence interval around this estimate, we next estimate the linear regression model in equation (3).

7.3 Regression Estimates of the Effect of Playing Partners' Ability on Own Score

The results of estimating equation (3) are shown in the first column of Table 3. Since our measure of ability is an average of varying numbers of prior adjusted scores, we weight all regressions by the number of past performance observations used to compute $Ability_i$. Shown in column (1), the coefficient on own ability is strongly statistically significant and large in magnitude, as expected. A one-stroke increase in a player's average score in past rounds is associated with an increase in that player's score of 0.678 strokes. That this coefficient is not equal to 1 suggests there is some measurement error in our measure of ability, but as a conditional reliability ratio this is reasonably large in magnitude. If we think of $\hat{\beta}_1$ this way, then as we described above $\hat{\gamma}_1/\hat{\beta}_1$ is a measurement-error corrected estimate of the effect of partners' ability on own score.

The estimate of γ_1 , the effect of playing partners' ability on own score, is not statistically significant, and the point estimate is actually negative. The insignificant point estimate suggests that improving the average ability of one's playing partners by one stroke actually increases (i.e.

worsens) one's score by 0.039 strokes. Our estimates make it possible to rule out positive peer effects larger than 0.045 strokes for an increase in average ability of one stroke. One stroke is about 50 percent more than one standard deviation in partners' average ability (0.66). If we divide the upper bound of the 95-percent confidence interval by the estimate of $\hat{\beta}$ to correct for measurement error, we can still rule out that a one-stroke increase in partners' average ability increases own score by more than 0.066 strokes. The results from our baseline specification therefore suggest that there are not significant peer effects overall. Column (2) verifies that the results are insensitive to controlling for player fixed effects instead of player ability. Going forward, we report results from the specification that includes own ability rather than player fixed effects for ease of interpretation.

7.4 The Effect of Different Dimensions of Ability: Does the Overall Effect Hide Evidence of Learning or Motivation?

As described earlier, we collected data on various dimensions of player skill. We hypothesize that players might learn about wind conditions or optimal strategies from more accurate players, or from better putters. In contrast, we assume that players cannot learn how to hit longer drives by playing alongside longer hitters. If these assumptions are correct, specifications comparable to (3) but replacing partners' average ability with partners' average driving distance, putts per round, or greens reached in regulation can separately identify learning and motivation effects. An effect of the accuracy measures would be interpreted as evidence of learning from peers, while an effect of partners' driving distance would be interpreted as evidence of motivation by peers.³²

The results are presented in columns (3)-(6) of Table 3. In column (3), we present results using average driving distance. While the coefficient on own driving distance is negative and strongly statistically significant (longer drives enable a player to achieve a lower score), the point estimate on partners' driving distance is small and statistically insignificant. The results for putts per round, shown in column (4), are similar. Own putting skill has a large and strongly significant effect on own score, but golfers do not appear to shoot lower scores when they play with better putters. The results for shot accuracy, shown in column (5), similarly show strong effects of own accuracy, but

³²One might argue that the former is a production complementarity in the typology set out at the outset of the paper, but the latter is clearly a purely social peer effect.

no effect of partners' accuracy on a golfer's performance.³³ Finally we present a specification that jointly estimates the effects of all three measures of ability in column (6). A golfer's putting accuracy and number of greens hit in regulation have the most significant effects on his score,³⁴ but as in the separately estimated specifications, no dimension of his partners' ability appears to have any effect on score.

7.5 Using Alternative Measures of Peer Ability

Having seen no evidence that the average ability of peers affects individual performance, we next ask whether the linear-in-means specification obscures peer effects in a different way. It is possible that it is not the mean ability of coworkers that matters, but rather the min or max. Possibly playing with bad players matters, but playing with good players does not. Or, maybe playing alongside one very good player or one very bad player affects performance. In each of these cases, the mean ability of peers would not measure the relevant peer environment accurately. Motivated by these possibilities, in Table 4 we present estimates of specification (3) where $\overline{Ability}_{-i,kt}$ is replaced with alternative measures of peers' ability. For comparison's sake, we report the base specification in column (1), which shows that peers' average ability has an insignificant effect on own score of -0.039. In column (2) we replace the average ability with the maximum ability of the player's peers. The point estimate is slightly smaller, but virtually unchanged. In column (3), we show that the estimated effect of the minimum of peers' ability is again negative and insignificant, and virtually the same as the average ability effect.

In columns (4) through (7), we investigate whether there appears to be a non-linear effect of partners' ability. To do this, we include indicators for whether individual i was assigned to a player in the top decile, top quartile, bottom quartile, or bottom quintile of the ability distribution in his category. None of the four estimates are statistically significant, though suggestively the point estimates for the top-quartile specifications are positive while those for the bottom-quartile

³³We also collected data on driving accuracy, specifically the fraction of fairways the golfer hits on tee shots. We do not include specifications using this variable because it does not strongly predict own score. This result is consistent with the work of Alexander and Kern (2005), who find large effects of putting accuracy on earnings, smaller effects of driving distance on earnings, and very small effects of driving accuracy on earnings. When we do estimate peer effects using driving accuracy, however, we find no effect of partners' driving accuracy on own score.

³⁴Two old golfing cliches seem to be consistent with the data: "Drive for show, putt for dough" and "Hit fairways and greens."

specifications are negative. Recall that lower scores are better, so this pattern would suggest that players play worse when they are matched with much better players. We also ask whether playing with Tiger Woods, the best player of his generation, affects performance. The point estimate suggests that being partnered with Tiger Woods reduces golfers' scores, but the standard errors are large enough that we cannot rule out a zero effect. In our sample, there are only 70 golfer-days paired with Tiger Woods.

The specifications thus far have assumed it is the absolute level of peers' ability that affects performance. An alternative hypothesis is that relative ability also matters. To investigate this possibility we present specifications that allow the effect of peers' ability to vary with the difference between peers' and own ability. These specifications are reported in columns (9) and (10) of Table 4. The results suggest that the effect of peers' ability does not vary with relative ability. Furthermore, there does not seem to be a non-linear relationship with relative ability. We have also looked at relative ability and non-linear effects using our other measures of ability and have found similarly small and statistically insignificant results.

7.6 Endogenous Effect Regressions and Common Shocks

Table 5 reports results of equation (4), the specification that replaces partners' ability with partners' score as the regressor of interest. As described earlier, the coefficient estimate on playing partners' score overstates the true peer effect if there are unobserved common shocks affecting all players uniformly in the group. Nevertheless, this regression is informative since it can be thought of as providing an upper bound on γ . The first column in Table 5 shows, contrary to the results above, that the peer effect is positive and statistically significant—an increase in the average score of one's playing partners is associated with an increase of own score by 0.071 strokes. An important point, however, is that without even accounting for the upward bias in this estimate due to common shocks, the coefficient is still small in magnitude. Mas and Moretti's (2006) elasticities evaluated at the mean of our dependent variable (own score), would predict that an increase in average partners' score of one stroke would raise own score by 0.170 strokes, more than two times what we estimate. We do not correct for measurement error here, since each player's score, and therefore the average

score of peers, is measured without error.³⁵

To look at the importance of common shocks more systematically, we try several additional controls. We hypothesize that the most likely sources of common shocks are variation in weather and crowd size. Because these shocks also affect the playing groups nearby on the golf course, we construct a control for common shocks which are the scores of the 8 groups ‘nearby’ on the golf course (i.e. the groups playing directly in front of and behind the reference group on the golf course).³⁶

The results in column (2) show that the introduction of nearby playing groups as a control reduces the point estimate on partners’ score by twenty percent (from 0.071 to 0.055), suggesting that common shocks play a role. The estimate on partners’ score, however, remains statistically significant. To verify that the controls for the performance of nearby groups is actually estimating a common shock we construct a placebo set of groups in the same tournament and verify that these placebo groups do not affect the point estimate on partners’ score. This estimate is reported in column (3). Indeed the estimate of γ_2 in columns (1) and (3) are virtually identical.

Lastly, we try two other controls for intraday common shocks. In column (4) we interact a time-of-day (early morning, mid-morning, afternoon) fixed effect with the full set of tournament fixed effects. This should capture weather shocks and other changes in course conditions that will affect all groups that play at the same part of the day (a common complaint on the PGA Tour is that afternoon groups experience more ‘spike marks’ on the green which make it more difficult to putt effectively). The point estimate on partners’ score drops again, and is almost identical to the estimate that includes neighboring groups score as a control, but the estimate of partners’ score is still statistically significant. To capture the fact that weather varies more smoothly than the dummy specification assumes, in column (5) we introduce a cubic in start-time which is allowed to vary

³⁵One might argue that score is a noisy measure of performance. Even though score is the measure by which players are judged, a golfer may play well but have a high score because of a few unlucky bounces. In this case, the measurement error correction used in the previous section will likely produce estimates that are too large because $Ability_i$ is an even noisier measure of the contemporaneous performance of player i than $Score_{-i}$ is of his partners’ contemporaneous performance.

³⁶In tournaments that use a 1-tee/10-tee split start, we also define nearby groups to be those groups that teed off at the same time but on a different part of the golf course. Those groups play the holes in a different order, but experience the same weather shock. We choose to use 8 nearby groups in order to get a more accurate estimate of the common shocks component. We conducted Monte Carlo simulations which demonstrated that including only one or two groups made it difficult to fully control for the common shocks when the variance of common shocks was larger than the variance in ability.

by tournament. The coefficient on partners' score is further reduced to 0.038 and becomes only marginally significant. In additional specifications, we include higher-order polynomials in start-time, which allow the effect of weather, and other common shocks that vary over the course of a day, to vary more and more and more flexibly. Moving across the columns, as the order of the polynomial increases from a cubic to a quartic to a quintic the estimated effect of partners' score decreases. With the control for a quintic in start-time, the endogenous peer effect coefficient is 0.022 and is insignificantly different from zero. Interestingly, adding controls for start-time or for the scores of neighboring groups does not appear to affect the estimate of own ability on own score. It appears that the correlations between own score and partners' score are driven primarily by common shocks.

As with any peer effects regression of own outcome on peer's outcome it is difficult to interpret these regressions. They should certainly be regarded as upper bounds for peer effects since any remaining common shocks that are not controlled for should bias the estimates upwards. Furthermore, regressions of outcomes on peers' outcomes suffer from the 'reflection problem' described by Manski (1993). In short, the fact that more and more extensive controls for common shocks reduce the estimate of γ_2 but do not appreciably affect the estimate of β_2 , along with the fact that the estimates of γ_1 are consistently zero, lead us to conclude that peer effects are negligible among professional golfers.

7.7 Do Peer Effects Vary with Player Skill or Tournament Prize Structure?

One possible explanation for why we find such different results than previous studies is that professional golfers are elite professionals subject to a selection process. There may exist heterogeneity in the susceptibility of workers to social influences by co-workers. And, perhaps the most successful professional golfers are those who are able to avoid these social responses. If heterogeneity in this ability across occupations explains the differences between our results and those in Mas and Moretti (2006), it may also be the case that there is heterogeneity among golfers in the susceptibility to social influences. In Table 6, we present estimates of equation (3) that allow the effect of partners' ability to vary by the reference player's skill. This interaction tells us, for example,

whether low-skill players respond more to high-skill players than high-skill players do.³⁷ The results are shown in columns (2) and (3). Although the coefficient is suggestively positive, it is not statistically significant at conventional levels. Note however, that the coefficient becomes larger and statistically significant if we drop category 1 players (see column (3)). A positive coefficient on this interaction term implies that lower-skilled players respond more to their co-workers' ability than better players do. The size of the standard errors prevents us from making strong conclusions, but the pattern of the point estimates in columns (2) and (3) is consistent with the idea that more skilled workers are less responsive to peer effects.³⁸

We next ask whether peer effects vary with the strength of financial incentives. In column (4) of Table 6 we present estimates from a specification where we allow the effect of playing partners' ability to vary with the total purse of the tournament.³⁹ A negative coefficient would imply that larger financial incentives reduce the peer effect. The estimated coefficient is negative but statistically insignificant. We have estimated specifications using measures of the dispersion of prizes (e.g. the ratio of the 90th and 10th percentile of earnings, the ratio of the top prize to the average prize, and the ratio of the standard deviation of prizes to the mean prize) and reproduce some of them in columns (4)-(8), but each specification produces large standard errors and inconsistent magnitudes. It is thus hard to conclude much from these results, but it is worth noting here that even in the tournaments in our sample with the smallest purses and the least convex prize structure, financial incentives are still very strong.

7.8 Why Are There No Peer Effects Among Professional Golfers?

Contrary to the findings from recent studies of peer effects in the workplace, we find no evidence that the ability or current performance of playing partners affects the performance of professional golfers. Mas and Moretti (2006) find large peer effects in a low-wage labor market where workers are not paid piece rates and do not have strong financial incentives to exert more effort. Bandiera,

³⁷It is worth pointing out that high- and low-skill are relative terms. All professional golfers are extremely high-skilled relative to the population.

³⁸The comparison of effects across categories would seem to imply no differential effect by ability. Recall, however, that there is a great deal of overlap of the ability distributions across categories. The highest ability category 3 players are actually very skilled, even relative to the average category 1 player. Thus, category is not a particularly good proxy for skill.

³⁹All purse values are adjusted to 2006 dollars using the PCE deflator.

Barankay and Rasul (2007) find peer effects specific to workers' friends in a low-skilled job where workers are paid piece rates. Experimental studies (e.g. Falk and Ichino, 2006) find evidence of peer effects in work-like tasks. In this section we speculate several explanations for our contrasting findings and conclude that there is much to be learned from the difference between our results and those of other recent studies.

Firstly, the PGA Tour is a unique labor market that is characterized by extremely large financial incentives for performance. In such a situation, it may be the case that the incentives for high effort are already so high that the marginal effect of social considerations are minimal, or zero. In similar labor markets where there are high-powered incentives for better performance (e.g. floor traders at an investment bank, lawyers in private practice, tenure track professors), the social effects of peers may not be as important as implied by existing studies. Consistent with this view is Mas and Moretti's (2006) conclusion that the peer effects they observe are mediated by co-worker monitoring. As incentives become stronger and monitoring output becomes easier, monitoring of effort becomes less necessary.

Perhaps just as interesting is the implication that social incentives may be a substitute for financial incentives. This would suggest that when creating strong financial incentives is difficult (such as when monitoring costs are high, or measuring individual output is difficult), firms should optimally organize workers to take advantage of social incentives.

Secondly, professional experience might lessen social influences like peer effects. The PGA Tour is a competitive workplace, and given the random assignment of peers, players who do not respond to the performance of their peers might be more consistently successful. It may be the case that as workers gain experience in highly-skilled occupations they learn how to avoid unproductive distractions or pressures not to make co-workers look bad.⁴⁰

Lastly and, we speculate, most importantly, the sample of workers under study has been subject to extreme selection. Many people play golf, but only the very best are professional golfers. Even among professionals, PGA Tour players are among the elite. It is quite possible that an important selection criterion is the ability to avoid the influences of playing partners. The results described at the end of the previous section are suggestively consistent with this view. Even among the highly

⁴⁰This is consistent with the experimental work of List and Millimet (2006) which demonstrates that market experience contributes to individual rationality.

selected group of professional golfers, the least skilled and least experienced golfers are the only ones whose productivity respond to the composition of their peers. We view this as an interesting conclusion because it suggests that there is a great deal of heterogeneity across individuals in their susceptibility to social influences in the workplace. It is an open question whether professional golfers are rare exceptions or representative of a larger class of skilled professional workers.

8 Conclusion

We use the random assignment of playing partners in professional golf tournaments to test for peer effects in the workplace. Contrary to recent evidence on supermarket checkout workers and soft-fruit pickers, we find no evidence that the ability or current performance of playing partners affects the performance of professional golfers. With a large panel data set we can observe players repeatedly, and the random assignment of players to groups makes it straightforward to estimate the causal effect of playing partners' ability on own performance. The design of professional golf tournaments also allows for direct examination of the role of common shocks, which typically make identification of endogenous peer effects difficult. In our preferred specification, we are able to reject positive peer effects of more than 0.045 strokes for a one stroke increase in playing partners' ability. We are also able to rule out that the peer effect is larger than 6.6 percent of the effect of own ability.

Interestingly, we find a positive effect of partners' score on own score, but we interpret this as mostly due to common shocks (and we present evidence that controlling for these common shocks reduces this correlation). The raw correlation in scores, though, might help explain why many PGA players perceive peer effects to be important.⁴¹ Our results suggest that players might be misinterpreting common shocks as peer effects.

We conclude by noting that though our results are different than those found in recent studies of peer effects in the workplace, we view our results as complementary. There is much to learn from the differences in findings. Primarily, our results suggest that there is heterogeneity in the

⁴¹For example, Darren Clarke (www.pga.com/pgachampionship/2004/news_interviews_081304_clarke.html) stated that it is easier to play better when everyone in the group is playing better because "you see good shots go into the green all the time and that makes it a lot easier to do the same yourself", and Billy Andrade (i.pga.com/pga/images/events/2006/pgachampionship/pdf/20060818_andrade.pdf) said "you kind of feed off each other and that's what we did."

importance of peer effects, both across individuals and across settings. We speculate that some people are better than others at avoiding the social effects found by Moretti and Mas (2006) and by Bandiera, Barankay and Rasul (2007). This ability may be one thing that characterizes workers in high-skilled labor markets. By focusing on low-skilled occupations, the existing studies miss this rich heterogeneity. Perhaps just as importantly, we show that peer effects are not important in a setting with strong financial incentives. This finding suggests that social effects may be substitutes for incentive pay, and is consistent at least in spirit with the work of List (2006) and Levitt and List (2007) who argue that the expression of social preferences is likely to vary according to whether behavior is observed in a market setting, the strength of incentives in that setting, and the selection of subjects that researchers observe. We hope our findings will spur other researchers to further explore this heterogeneity in peer effects in the workplace.

References

- [1] Alexander, Donald and William Kern (2005). "Drive for Show and Putt for Dough?: An Analysis of the Earnings of PGA Tour Golfers," *Journal of Sports Economics*, 6: 46-60.
- [2] Andreoni, Jim and Doug Bernheim (2006). "Social Image and the 50-50 Norm," Working Paper.
- [3] Bandiera, Oriana, Iwan Barankay and Imran Rasul (2007). "Social Incentives in the Workplace," University College London Mimeo. July.
- [4] Bandiera, Oriana, Iwan Barankay and Imran Rasul (2005). "Social Preferences and the Response to Incentives: Evidence from Personnel Data," *Quarterly Journal of Economics* 120(3): 917-962.
- [5] Benabou, Roland and Jean Tirole (2002). "Self-Confidence and Personal Motivation," *Quarterly Journal of Economics*, 117(3): 871-915.
- [6] Bertrand, Marianne, Erzo F.P. Luttmer and Sendhil Mullainathan (2000). "Network Effects and Network Cultures," *Quarterly Journal of Economics*, 115: 1019-1055.
- [7] Bronars, Stephen and Gerald Oettinger (2000). "Effort, Risk-Taking, and Participation in Tournaments: Evidence from Professional Golf," University of Texas Mimeo.
- [8] Bognanno, Michael and Ronald Ehrenberg (1990). "Do Tournaments have Incentive Effects?" *Journal of Political Economy* 98: 1307-1324.
- [9] Boozer, Michael and Stephen Cacciola (2001). "Inside the 'Black Box' of Project Star: Estimation of Peer Effects using Experimental Data." Working Paper.
- [10] Falk, Armin and Andrea Ichino (2006). "Clean Evidence on Peer Effects," *Journal of Labor Economics* 24(1): 39-57 .
- [11] Festinger, Leon (1954). "A theory of social comparison processes," *Human Relations*: 117-140.

- [12] Glaeser, Edward L., Bruce Sacerdote and Jose A. Scheinkman (1996). "Crime and Social Interactions," *Quarterly Journal of Economics* 111: 507-548.
- [13] Graham, Bryan (2007). "Identifying Social Interactions through Excess Variance Contrasts," University of California, Berkeley Working Paper.
- [14] Kroft, Kory (forthcoming, 2007). "Takeup, Social Multipliers and Optimal Social Insurance," *Journal of Public Economics*.
- [15] Lazear, Edward, Ulrike Malmendier and Roberto Weber (2006). "Sorting in Experiments with Application to Social Preferences"
- [16] Lemieux, Thomas, Bentley Macleod and Daniel Parent (2006). "Performance Pay And Wage Inequality" Departmental Working Papers, McGill University, Department of Economics.
- [17] Levitt, Steven and John List (2007). "What Do Laboratory Experiments Measuring Social Preferences Tell Us About the Real World?" *Journal of Economic Perspectives*, Forthcoming.
- [18] List, John A. (2006). "The Behavioralist Meets the Market: Measuring Social Preferences and Reputation Effects in Actual Transactions," *Journal of Political Economy*, 114(1): 1-37.
- [19] List, John A. and Daniel L. Millimet (2006). "Bounding the Impact of Market Experience on Rationality: Evidence from a Field Experiment with Imperfect Compliance," University of Chicago mimeo.
- [20] Moffitt, Robert (2001). "Policy Interventions, Low-Level Equilibria, and Social Interactions." In S. Durlauf and P. Young (Eds.), *Social Dynamics*. Cambridge: MIT Press.
- [21] Manski, Charles (1993). "Identification of Endogenous Social Effects: The Reflection Problem," *Review of Economic Studies* 60: 531-542.
- [22] Mas, Alexandre and Enrico Moretti (2006). "Peers at Work," NBER Working Paper #12508.
- [23] Orszag, Jonathan (1994). "A New Look at Incentive Effects in Golf Tournaments," *Economic Letters* 46: 77-88.

- [24] Pessin, J. (1933). "The comparative effects of social and mechanical stimulation on memorizing," *American Journal of Psychology* 45: 263–270.
- [25] Sacerdote, Bruce (2001). "Peer Effects with Random Assignment: Results for Dartmouth Roommates." *Quarterly Journal of Economics*, 116(2): 681-704.
- [26] Saez, Emmanuel and Esther Duflo (2003). "The Role of Information and Social Interactions in Retirement Plan Decisions: Evidence From a Randomized Experiment," *Quarterly Journal of Economics* 118: 815-842.
- [27] Triplett, Norman (1898). "The Dynamogenic Factors in Pacemaking and Competition," *American Journal of Psychology* 9: 507-533.
- [28] Zajonc, R. B (1965). "Social facilitation," *Science* 149: 269–274.
- [29] Zimmerman, David J. (2003). "Peer Effects in Higher Education: Evidence from a Natural Experiment, Williams Project on the Economics of Higher Education," *Review of Economics and Statistics*, February, 85(1): 9-23.

Appendix A: Introduction to Golf

The objective in golf is to take the fewest number of strokes on each hole to hit a 1.7-inch golf ball into a 4.25-inch cup. The player with the lowest total number of strokes over all holes is the winner. Golf holes typically are between 150 and 600 yards in length, and they often have ‘hazards’ (e.g. sand traps, water hazards, out-of-bounds areas, and tall trees) that require skill and strategy to avoid. The cup is placed on an area of shorter grass called the ‘green’. The grass is cut shorter on the green so that the ball rolls smoothly when it is struck. Such rolling shots are called ‘putts’.

The first shot on each hole is the tee shot. All players hit their tee shot from the same tee area. This is the only shot where the player is allowed to place the ball on a wooden tee which holds the ball slightly above the ground. There is an area between the tee area and the green called the ‘fairway’ where the grass is cut short (though longer than the grass on the green). On either side of the fairway is longer grass called the ‘rough’. It is easier to hit the ball when it is sitting on shorter grass.

There are three types of holes: par-3 holes, par-4 holes, and par-5 holes. The par indicates the number of strokes in which the hole is expected to be completed. On a par-3 hole, the typical strategy is to attempt to hit the tee shot directly onto the green. For a par-4 hole, the typical strategy is to reach the green in two shots. For a par-5 hole, the typical strategy is to reach the green in three shots.⁴²

Inside the cup is a 7 foot flagstick called the ‘pin’. The pin allows the players to see from a distance where the cup is on the putting green. Often the cup is placed near the edge of the green and near hazards (e.g. sand traps and water hazards), creating a risk-reward trade-off: if a player wants to try to hit his approach shot near the cup, he increases the chance that he hits into a hazard.

Many golf shots present similar risk-reward trade-offs. For example, on a par-4 a player might want to hit a driver (which is the golf club that allows the player to hit the ball the farthest, but which is the most difficult to control) to leave himself the shortest approach shot into the putting green (since players have more control and accuracy over shorter approach shots). However, it is

⁴²A combination of better golf technology and more athletic golfers, however, have transformed most par-5 holes on the PGA Tour into ‘long par 4-holes’ where a majority of the golfers can reach the putting green in one approach shot.

harder to hit the tee shot onto the fairway with a driver, and tee shots that miss the fairway and land in the rough leave more difficult approach shots.

Appendix B: Player Categories and Random Assignment

Based on their past performance over their career, players are placed in one of four categories: 1, 1A, 2, and 3. The categories are assigned using the following rules:

- Category 1 players are tournament winners and the top 25 money winners from the previous year and PGA Tour life members (e.g. Tiger Woods, Phil Mickelson and Ernie Els).
- Category 1A players include former champions of the four majors and The Players Championship (e.g. former British Open champions John Daly and Nick Faldo).
- Category 2 players include those in the top 125, players with 50 or more career cuts made and players in the top 50 of the World Golf Rankings. This encompasses the majority of the field most weeks.
- Category 3 includes all others, such as local qualifiers. These golfers get the first and last tee times of each session. When Annika Sorenstam became the first woman in 58 years to play in a PGA Tour event, she was a Category 3 player and was assigned the earliest and latest tee times in her two rounds.

Although there are four categories of players, the 1A and 2 players are paired together randomly, so for our purposes we re-label 1A players as category 2 players and work only with three categories of players. There is an additional complication, though, which is that players might be paired with players from a different category if the number of players in a given category is not a multiple of three. In that case category 1A players are paired with category 1 players and category 3 players are paired with category 2 players.

Categories are assigned at the beginning of the season, and for the most part the category assignments are static. However, if a player wins a tournament or enters the top 25 money list, then that player can be ‘promoted’ to Category 1 status during the season. Likewise, if a player drops out of the top 50 World Golf Ranking and does not satisfy the other definitions of a Category 2 player, then that player is ‘demoted’ to Category 3 status during the season. We are never able to directly observe players changing categories in our data. In the 2002 season, we have two ‘snapshots’ of category status directly from the PGA Tour (at the beginning of the season and halfway through the season). In the 2005-2006 seasons, we were not able to get the category status of the players on

the tour, but we used the definitions above to assign players as best we could at the beginning of the season, and then we used a probabilistic matching algorithm to assign the remaining players. We tested the matching algorithm on the 2002 season (where we had the categories given to us directly from the PGA Tour) and we verified that we got more than 99 percent of category assignments correct.

The algorithm works as follows: we start with a list of players where we are sure we know the category status throughout the season (this is most obvious for elite players, former major champions, and former Nationwide Tour players). Then we look at every playing partner of those players during the season and assign the playing partners to the same category. For players who get matched to different categories, we flag them and manually decide which category they belong to.

Using these categories we test for random assignment by tournament and we drop the following tournaments which fail the test for random assignment: The Masters, U.S. Open, British Open, PGA Championship, Walt Disney Championship, Tour Championship, Players Championship.

Figure 1: Tournament Earnings

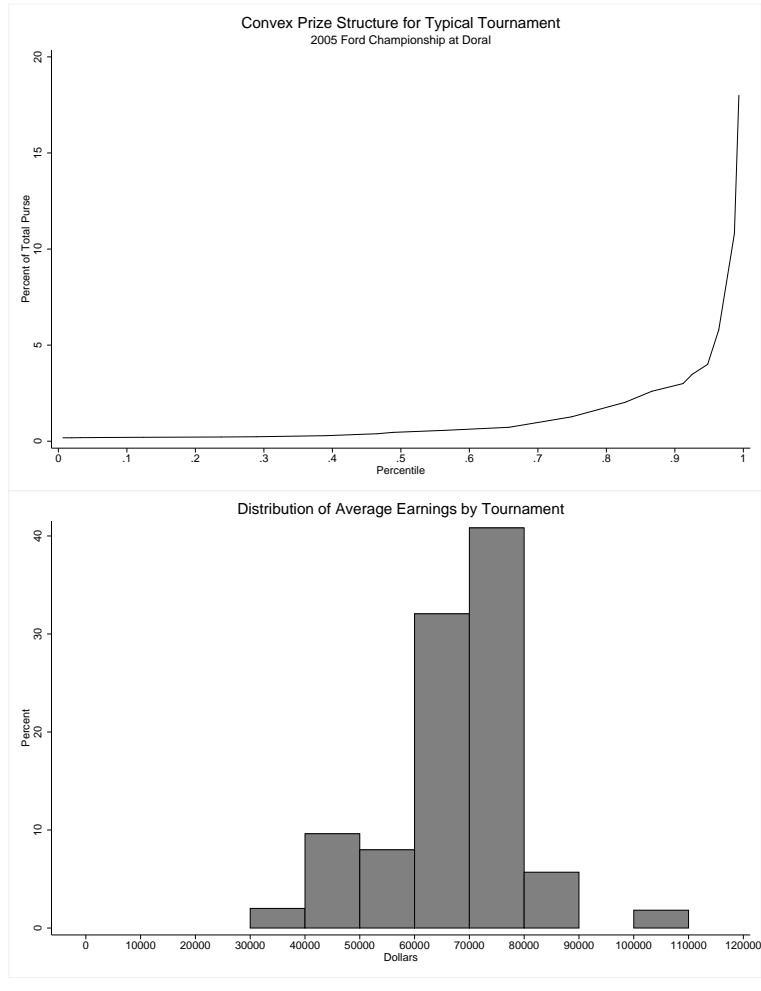


Figure 2: Monte Carlo of OLS randomization tests

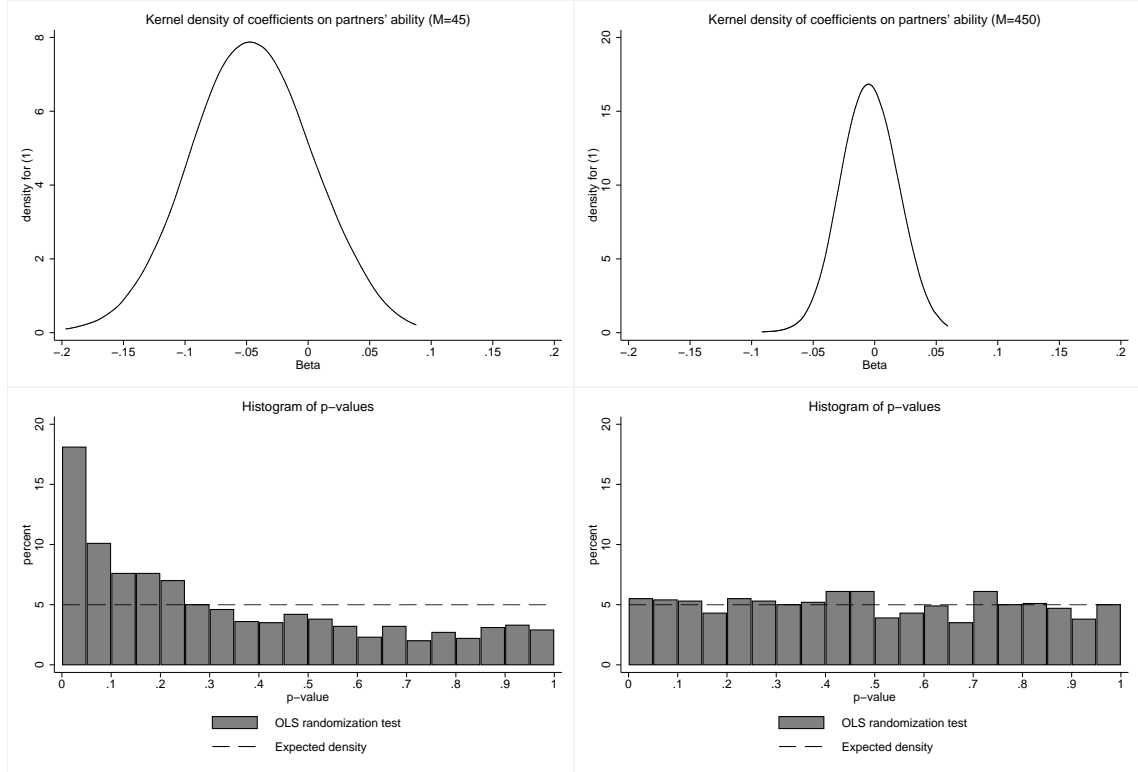


Figure 3: Monte Carlo of modified randomization tests

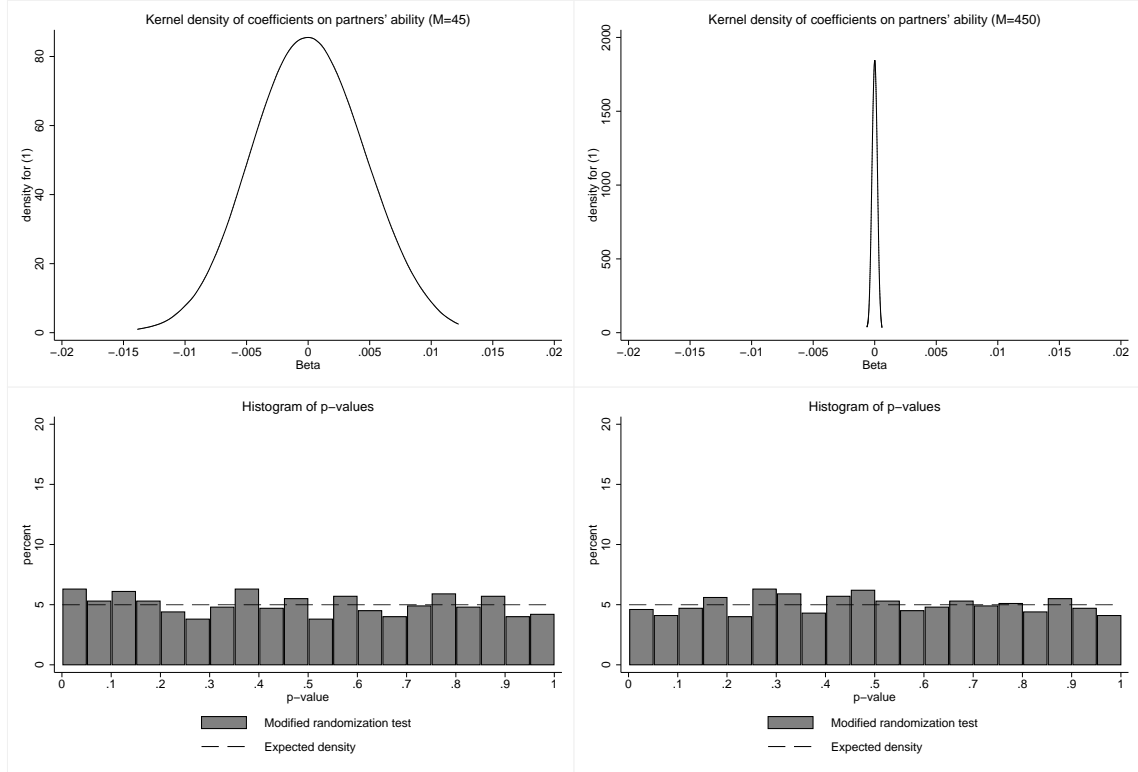


Figure 4: Kernel density of player handicaps

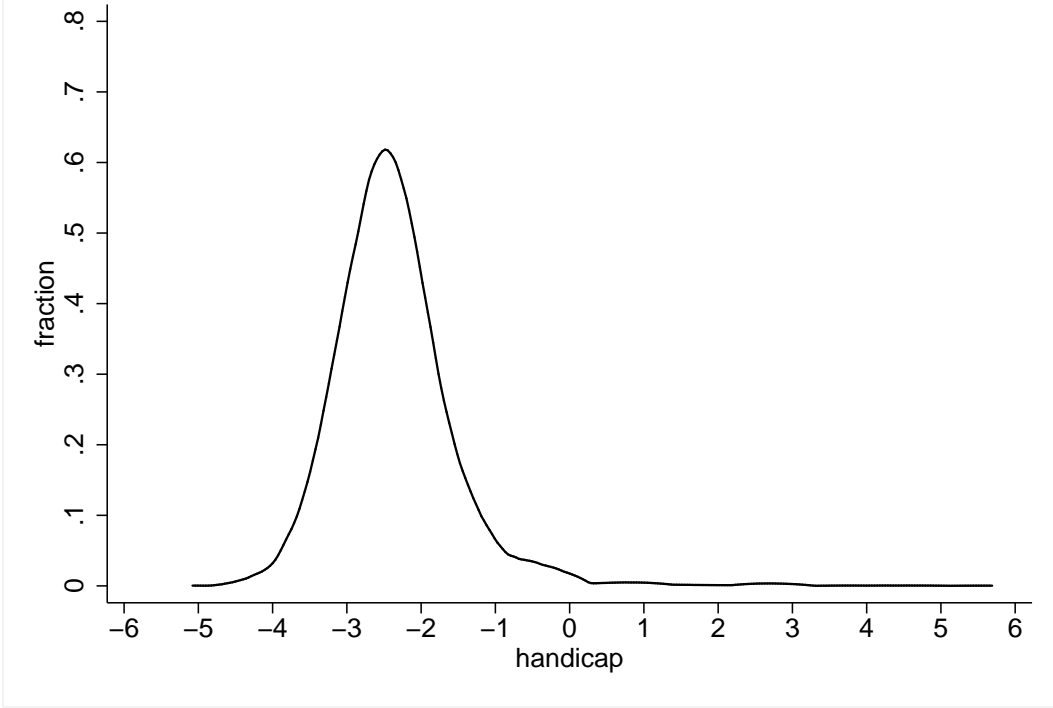


Figure 5: Kernel density of player handicaps, by category

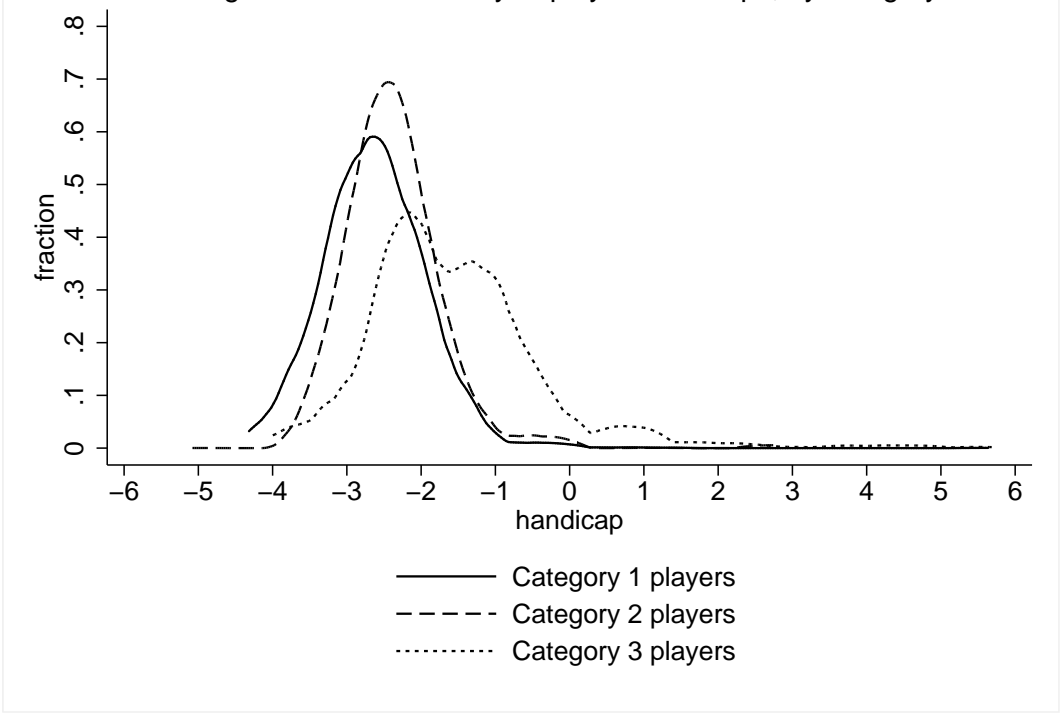


Figure 6: Scores for the entire sample

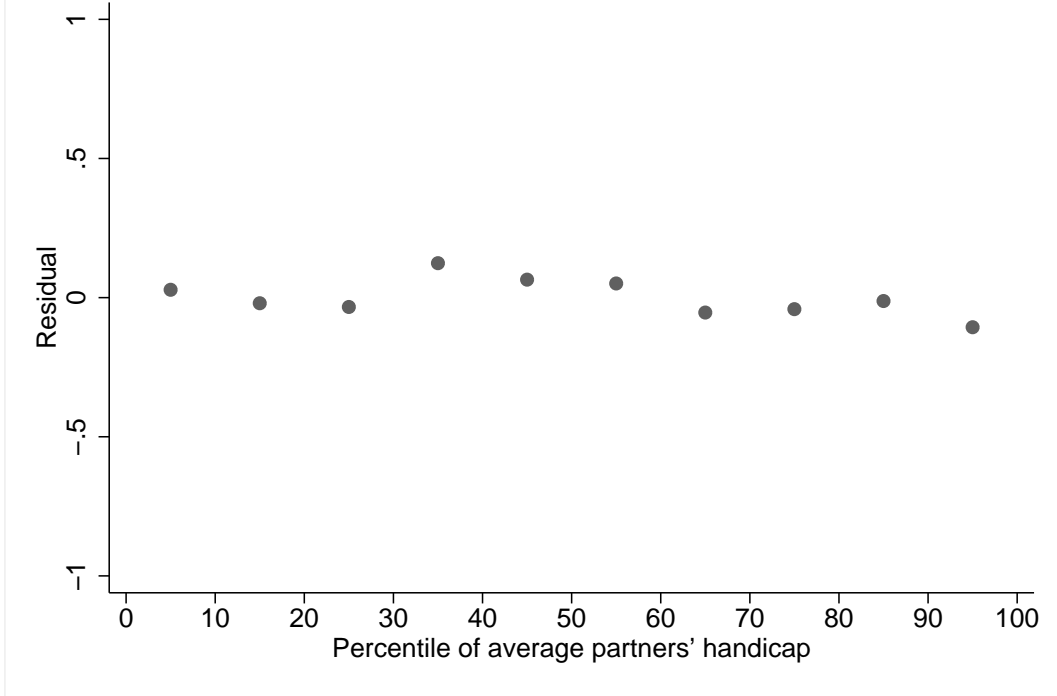


Table 1
Verification of Random Assignment

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Avg(Ability), partners	0.159 (0.021)	0.116 (0.020)	-0.019 (0.016)				-0.017 (0.015)	-0.004 (0.017)	-0.063 (0.046)
Avg(Driv. Dist.), partners				0.007 (0.014)					
Avg(Putts), partners					-0.003 (0.012)				
Avg(Greens/rd), partners						-0.001 (0.014)			
Tourn F.E.	Y	N	Y	Y	Y	Y	Y	Y	Y
Cat F.E.	N	Y	Y	Y	Y	Y	Y	Y	Y
Tourn*Cat F.E.	N	N	Y	Y	Y	Y	Y	Y	Y
Keep only Cat 1	N	N	N	N	N	N	Y	N	N
Keep only Cat 2	N	N	N	N	N	N	N	Y	N
Keep only Cat 3	N	N	N	N	N	N	N	N	Y
R ²	0.133	0.158	0.626	0.727	0.545	0.723	0.851	0.735	0.595
N	8750	8750	8750	8587	8587	8587	2947	5067	736

Notes:

- a. Results from estimating equation (2).
- b. Avg(Ability) is the average handicap of a player's playing partners, as described in the text.
- c. All specifications include the average of all of the other players in your tournament-by-category urn (not including yourself). This control is necessary to produce a well-behaved randomization test.
- d. All specifications include tournament fixed effects. Note that the 'urn' in column (1) is defined to be tournament cells instead of tournament-by-category cells (this confirms that our test can detect non-random assignment).
- e. Columns (3)-(9) include tournament-by-player-category interaction dummies to pick up selection into tournaments.
- f. Columns (7)-(9) impose various sample restrictions.
- g. Standard errors are in parentheses and are clustered by playing group.
- h. All specifications are unweighted (unlike the specifications in the other tables), since the weighting is not appropriate for the randomization test.

Table 2
Descriptive Statistics

<u>Variable</u>	<u>Obs</u>	<u>Mean</u>	<u>Std Dev.</u>	<u>Min</u>	<u>10th</u>	<u>percentiles</u>				<u>Max</u>
						<u>25th</u>	<u>50th</u>	<u>75th</u>	<u>90th</u>	
Score	17403	71.182	3.183	61	67.0	69.0	71.0	73.0	75.0	89
Avg(Score), partners	17403	71.236	2.490	62	68.0	69.5	71.0	73.0	74.5	85
Ability (Handicap)	17403	-2.38	0.78	-4.32	-3.19	-2.83	-2.45	-2.03	-1.54	5.635
Avg(Ability), partners	17403	-2.37	0.66	-5.08	-3.05	-2.77	-2.44	-2.10	-1.68	5.285
Driving distance (yards)	17077	281.02	9.88	238.80	269.10	274.04	280.10	288.00	293.30	324
Putts per round	17077	28.67	1.03	24.07	27.34	28.04	28.83	29.34	29.72	33.500
Greens per round	17077	11.60	0.65	7.00	10.86	11.33	11.66	11.97	12.32	14.994
Purse (in millions, 2006\$)	15985	4.40	0.85	2.27	3.24	3.78	4.32	5.11	5.50	6.000
log(Purse)	15985	15.28	0.20	14.63	14.99	15.14	15.28	15.45	15.52	16
Tiger Woods in group	17403	0.00	0.06	0.00	0.00	0.00	0.00	0.00	0.00	1
<u>Category 1 players</u>										
Ability	5863	-2.62	0.68	-4.32	-3.46	-3.05	-2.69	-2.22	-1.75	5.635
Avg(Ability), partners	5863	-2.62	0.50	-4.25	-3.23	-2.95	-2.63	-2.31	-2.00	0.562
<u>Category 2 players</u>										
Ability	10075	-2.37	0.64	-3.57	-3.10	-2.75	-2.44	-2.05	-1.72	2.741
Avg(Ability), partners	10075	-2.36	0.53	-5.08	-2.94	-2.70	-2.40	-2.10	-1.79	2.741
<u>Category 3 players</u>										
Ability	1465	-1.46	1.20	-4.00	-2.62	-2.19	-1.54	-0.94	-0.24	5.285
Avg(Ability), partners	1465	-1.49	1.10	-4.00	-2.62	-2.21	-1.60	-0.99	-0.28	5.285

Notes:

a. See Appendix B for more information on player categories.

Table 3
The Effect of Peers' Ability on Own Score

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Own Ability	0.678 (0.084)						0.949 (0.074)	0.544 (0.077)	0.784 (0.129)	0.566 (0.071)
Avg(Ability), partners	-0.039 (0.042)	-0.026 (0.040)					0.013 (0.095)	-0.068 (0.071)	-0.003 (0.115)	-0.057 (0.062)
Driving Distance			-0.028 (0.007)			-0.015 (0.007)				
Avg(Driving), partners			0.002 (0.005)			0.002 (0.005)				
Putts				0.144 (0.053)		0.161 (0.053)				
Avg(Putts), partners				-0.031 (0.040)		-0.036 (0.040)				
Greens per round					-0.636 (0.080)	-0.595 (0.084)				
Avg(Greens), partners					-0.004 (0.063)	-0.005 (0.065)				
Tourn*Cat F.E.	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Player F.E.	N	Y	N	N	N	N	N	N	N	N
Cat=1 only	N	N	N	N	N	N	Y	N	N	N
Cat=2 only	N	N	N	N	N	N	N	Y	N	N
Cat=3 only	N	N	N	N	N	N	N	N	Y	N
Drop Cat 1	N	N	N	N	N	N	N	N	N	Y
R ²	0.143	0.208	0.131	0.130	0.138	0.141	0.152	0.122	0.211	0.129
N	17403	17403	17077	17077	17077	17077	5863	10075	1465	11540

Notes:

- a. Results from baseline specifications as specified in equation (1).
- b. The dependent variable is the golf score for the round.
- c. The Ability variable is measured using the player's handicap, and "Avg(Ability), partners" is measured using the average of your playing partners' handicaps. The Avg() operator takes the average of the variable for for the playing partners in your group.
- d. All specifications include tournament fixed effects, player category fixed effects, and tournament-by-player-category interaction dummies. Also, all specifications include a control for the mean ability of the other players in your tournament-by-category urn (where the mean is calculated over all other players in your cell not including yourself).
- e. Standard errors are in parentheses and are clustered by playing group.
- f. All specifications weight each observation by the number of tournaments used to construct the player's handicap.

Table 4
Peer Effects with Alternative Measures of Peer Ability

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Own Ability	0.678 (0.084)	0.678 (0.077)	0.677 (0.077)	0.678 (0.077)	0.676 (0.077)	0.678 (0.077)	0.680 (0.077)	0.679 (0.077)	0.674 (0.079)	0.681 (0.080)
Avg(Ability), partners		-0.039 (0.042)							-0.035 (0.062)	-0.040 (0.062)
Max(Ability), partners		-0.023 (0.039)								
Min(Ability), partners)			-0.038 (0.052)							
1{any partner in top 10%}				0.036 (0.072)						
1{any partner in top 25%}					0.061 (0.064)					
1{any partner in bot 25%}						-0.022 (0.075)				
1{any partner in bot 10%}							-0.102 (0.157)			
Tiger Woods is partner								-0.354 (0.500)		
Avg(Ability), partners *									-0.004 (0.021)	-0.035 (0.047)
(Avg(Ability), partners - Own Ability)										
Avg(Ability), partners *										0.008 (0.009)
(Avg(Ability), partners - Own Ability) ²										
Tourn*Cat F.E.	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
R ²	0.143	0.143	0.143	0.143	0.143	0.143	0.143	0.143	0.141	0.141
N	17403	17370	17370	17396	17396	17396	17396	17403	17403	17403

Notes:

- a. Column (1) is reproduced from Table 3. Other columns present results from modifying baseline specifications as specified in equation (1) to support heterogeneous peer effects.
- b. The dependent variable is the golf score for the round.
- c. The Ability variable is measured using the player's handicap, and "Avg(Ability), partners" is measured using the average playing partners' handicap. The Avg() operator takes the average of the variable for the playing partners in your group.
- d. All specifications include tournament fixed effects, player category fixed effects, and tournament-by-player-category interaction dummies. Also, all specifications include a control for the mean ability of the other players in your tournament-by-category urn (where the mean is calculated over all other players in your cell not including yourself).
- e. Standard errors are in parentheses and are clustered by playing group.
- f. All specifications weight each observation by the number of tournaments used to construct the player's handicap.

Table 5
The Effect of Peers' Score on Own Score

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Own Ability	0.682 (0.074)	0.651 (0.094)	0.681 (0.077)	0.684 (0.076)	0.680 (0.076)	0.642 (0.095)	0.677 (0.076)	0.683 (0.077)
Avg(Score), partners	0.071 (0.012)	0.055 (0.019)	0.072 (0.016)	0.056 (0.017)	0.038 (0.017)	0.022 (0.019)	0.030 (0.017)	0.022 (0.017)
Avg(Score), neighboring groups		0.159 (0.015)				0.130 (0.017)		
Avg(Score), placebo groups			-0.020 (0.019)					
Tourn*Cat F.E.	Y	Y	Y	Y	Y	Y	Y	Y
Tourn*Time-of-day F.E.	N	N	N	Y	N	N	N	N
Time cubic per tourney	N	N	N	N	Y	Y	N	N
Time quartic per tourney	N	N	N	N	N	N	Y	N
Time quintic per tourney	N	N	N	N	N	N	N	Y
R ²	0.143	0.159	0.143	0.151	0.161	0.18	0.165	0.17
N	17403	13171	17403	17403	17403	13171	17403	17403

Notes:

- a. Results from alternate specifications using average partners' score instead of average partners' ability as the primary independent variable of interest.
- b. The dependent variable is the golf score for the round.
- c. The Ability variable is measured using the player's handicap, and "Avg(score), partners" is the average of the score of the playing partners. The placebo groups are formed by drawing a random sample of groups in the same tournament.
- d. All specifications include tournament fixed effects and player category fixed effects.
- e. All columns include tournament-by-player-category interaction dummies to pick up selection into tournaments.
- f. Standard errors are in parentheses and are clustered by playing group.
- g. All specifications weight each observation by the number of tournaments used to construct the handicap.

Table 6
Do Peer Effects Vary with Player Skill or Tournament Prize Structure?

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Own Ability	0.678 (0.084)	0.752 (0.052)	0.538 (0.071)	0.702 (0.080)	0.702 (0.080)	0.349 (0.158)	0.697 (0.080)	0.697 (0.080)	0.352 (0.158)
Avg(Ability), partners	-0.039 (0.042)	-0.041 (0.052)	-0.083 (0.062)	-0.055 (0.054)	-0.055 (0.054)	-0.062 (0.076)	-0.054 (0.054)	-0.054 (0.054)	-0.062 (0.076)
Avg(Ability), partners * Own Ability		0.087 0.065	0.191 (0.075)						
Avg(Ability), partners * Purse				-0.027 (0.063)			-0.034 0.063		
Avg(Ability), partners * log(Purse)					-0.114 (0.255)			-0.144 (0.258)	
Avg(Ability), partners * (90-10 prize ratio)						0.114 (0.096)			0.111 (0.097)
Own Ability * Purse							-0.104 (0.060)		
Own Ability * log(Purse)								-0.104 (0.060)	
Own Ability * (90-10 prize ratio)									-0.120 (0.108)
Tourn*Cat F.E.	Y	Y	Y	Y	Y	Y	Y	Y	Y
Drop Cat 1	N	N	Y	N	N	N	N	N	N
R ²	0.143	0.143	0.129	0.142	0.142	0.179	0.143	0.143	0.179
N	17403	17403	11540	15985	15985	5824	15985	15985	5824

Notes:

- a. Column (1) is reproduced from Table 3. Other columns present results from modifying baseline specifications as specified in equation (1) to support heterogeneous peer effects.
- b. The dependent variable is the golf score for the round.
- c. The Ability variable is measured using the player's handicap, and Avg(partners' Ability) is measured using the average playing partners' handicap. The Avg() operator takes the average of the variable for for the playing partners in your group.
- d. All specifications include tournament fixed effects, player category fixed effects, and tournament-by-player-category interaction dummies. Also, all specifications include a control for the mean ability of the other players in your tournament-by-category urn (where the mean is calculated over all other players in your cell not including yourself).
- e. Standard errors are in parentheses and are clustered by playing group.
- f. All specifications weight each observation by the number of tournaments used to construct the player's handicap.