Canadian Labour Economics Forum

WORKING PAPER SERIES

Identifying the Value of Teamwork: Application to Professional Tennis

Kevin Devereux (University of Toronto)

Fall 2018, WP #14

Identifying the Value of Teamwork: Application to Professional Tennis*

Kevin Devereux†

(Click here for current version)

September 12, 2018

Abstract

Do workers vary in their ability to work with others? I compare a given worker's productivity in solitary production to their value-added to team production to identify *team skills*: a worker's contribution to team production above and beyond that given by general skills. The identifying assumption is that workers use general skills in both production functions, but team skills only in team production. Professional men's tennis provides a useful setting to compare solo work (singles) to teamwork (doubles). I find that around 50% of variation in team output is explained by team skills. This is robust to a variety of specifications, including nonlinearities in player inputs. Players sort positively-assortatively along both skill dimensions, yielding indirect returns to skills of about half the magnitude of the direct returns.

JEL codes: C33, D31, J24, J31

Keywords: skills, human capital, teamwork, sorting, non-partite matching, assortative matching.

†Thanks to Aloysius Siow, Laura Turner, Ronald Wolthoff, Louis-Philippe Morin, Angelo Melino, Jiaying Gu, David Walker-Jones, Hammad Shaikh, and the seminar participants at the University of Toronto for helpful discussions and feedback. Remaining errors are my own. Contact: Department of Economics, University of Toronto, 150 St. George St., Toronto, Ontario M5S 3G7, Canada. Email: kevin.devereux@ucd.ie.

1 Introduction

Virtually all economic output is produced by teams. Employers and analysts in a broad range of professions emphasize the importance of teamwork, and economists have found that participation in team activities predicts future labour market returns.[1] Teamwork is growing: since the 1980s the share of workers at large US firms working in teams has increased, along with investment in team-building (Lazear and Shaw 2007).

Crucially, the best 'team players' may not be the most productive when working alone.[2] If workers exhibit heterogeneity in the types of skills required for effective teamwork, some workers have a *comparative advantage* at teamwork that cannot be predicted by solitary productivity measurements. This implies that a worker's human capital cannot be effectively summarized by a single dimension when considering allocation across teams of varying sizes.

This paper compares a given worker, performing a given job, alone and as part of a team. I exploit the labour market structure of men's professional tennis to compare productivity in exogenously set team sizes of one (singles tennis) and two (doubles tennis). From a worker's solitary productivity I create an index of *general skill*. By observing mobility across doubles partners, I identify a player's value-added to a team using the two-way fixed effects model due to Abowd et al. (1999).[3] I define *team skill* as the portion of composite skill unexplained by general skill: the systematic contribution of a worker to team output that is not predicted by solitary performance.

I find large and significant variation in team skill across players. Team skill is a key determinant of productivity, explaining around 50% of across-team output variation (compared to under 20% due to general skill). This is not an artifact of nonlinearity

[1] Kuhn and Weinberger (2005) emphasize the association between performing leadership roles in high school – such as captaining a sports team or chairing a school club – and higher wages years after graduation. A reported but unemphasized result is that the effect of team participation is twice as large as team leadership. Stevenson (2010) establishes a causal effect of team sports participation on wages based on variation in Title IX rollout across states, supporting the hypothesis that "players are taught to function as a team. The development of these skills could be especially important for girls who must try to maneuver their way through traditionally male occupations later in life" (Stevenson 2010, p. 287).

[2] Arcidiacono et al. (2017) find disjoint variation between a worker's own productivity compared and their ability to increase the productivity of teammates in an application to professional basketball. In contrast to the current paper, all of their productivity measures are taken from team production (see Appendix B for details).

[3] Intuitively: the common component of the teams ij and ik is worker i; if both teams perform well, the model attributes this to i's skills. But if ij perform well and ik poorly, the model identifies k as the high-skilled worker.

in players' inputs into team production – I conduct a specification test that fails to reject the two-way fixed effects model.[4] Nor is it due to idiosyncratic productivity of certain teams – the data fail to reject the two-way fixed effects model in favour of the more general (and much more flexible) match effects model.[5] Reduced form evidence suggests that potential partners cannot anticipate idiosyncratic match quality before playing together, validating the assumptions of the two-way fixed effects model.

These results indicate that solitary productivity is a poor predictor of team productivity in the current setting, but that a worker's productivity on one team is a good predictor of his productivity on another. Moreover, some workers have a comparative advantage at teamwork compared to solitary work, and so should optimally be assigned to group tasks, while others should work alone.

A second set of results concerns sorting. Players sort positively-assortatively by general skills as well as by team skills, but do not sort systematically across the two dimensions.[6] This implies that a given worker's returns to skills are mediated by sorting, with higher skilled players benefiting from being matched to higher skilled partners in addition to receiving direct returns to skills. Indirect returns are in the order of half the size of direct returns, with a one standard deviation increase in general skill increasing doubles team productivity by 38% directly and 20% by virtue of being matched to a better partner on average. The direct and indirect returns to a one standard deviation increase in team skill are 66% and 35% respectively. Because prize money is paid based on productivity in professional tennis, these figures translate directly into higher earnings.[7]

To my knowledge this is the first paper to separately identify general and team skills. While it is widely understood that workers can influence the performance of their teammates, the large literature on peer effects implicitly assumes that it is the best individual workers who have the highest spillovers onto team members – an assumption

[4] Moreover, the nonlinear extension yields similar quantitative results to the two-way fixed effects model. See Appendix D for details.

[5] The 'saturated' match effects model allows unrestricted variation in productivity across teams.

[6] A high general skilled player is likely to be matched to a high general skilled partner, and a high team skilled player to a high team skilled partner; but a high general skilled player is no more likely than chance to be matched to a high team skilled partner.

[7] This may not be the case when earnings are determined by contracting rather than a predetermined prize money schedule. Arcidiacono et al. (2017) find that NBA players are compensated for their own propensity to score, but not for the effect their presence has on the scoring propensity of teammates. They note that parsing individual contributions to teamwork can be particularly difficult. Sarsons (2017) finds that female academic economists are rewarded for solo authored publications, but not for coauthored publications (while men are rewarded for both).

that is rejected in the current setting.

Separately identifying the general and team skill components of a worker's contribution to productivity is difficult, especially if output is only observed at the team level. Past literature has exploited information on intermediate outcomes: byproducts of teamwork that contribute to team output. To make meaningful inference about specific team members, these outcomes must be associated with an individual. Arcidiacono et al. (2017) separately identify a worker's own skill and a spillover ability by using information on identity of the scorer in professional basketball. They infer a player's 'own skill' from his likelihood of scoring during a given possession, and 'spillover skill' from the effect of his presence on the likelihood of a teammate scoring. The identification assumption is that own skill causes own scoring but not scoring by others; and that spillover skill causes others to score but not oneself.

The current paper makes a qualitatively different (and novel) identifying assumption: workers always employ general skills, but employ team skills only when working with others. This defines team skill as a worker's contribution to team production net of solitary ability. Both modes of production are quantified in terms of output, allowing both skill classes to be measured in units of value-added.

Professional tennis provides a useful setting to apply this identification technique. Players participate in singles and doubles categories with the same objective: get the ball over the net. A player's differential ability to apply skills in the team setting forms the basis of the team skill measurement.

The key advantage of the Association of Tennis Professionals (ATP) data is the observation individual production. Occurring outside the context of any team, singles tennis provides a clean signal of individual productivity uncontaminated by the influence of teammates.[8] In contrast, any metric determined within the team setting – even one associated with an individual – may be influenced by teammates. Using intermediate outcomes from team production as individual productivity signals thus requires functional form assumptions on how the inputs of different team members influence the outcome measurement associated with a given worker. If outcomes from the team setting are endogenous to a team-level optimization problem that allocates different workers into different roles that determine observable outcomes, workers who vary within a single skill dimension may produce differential influence on teammates. This speaks to a

[8] Typical individual productivity signals, such as grade point average, IQ score, and tests on aptitude exams may be independent of a partner's influence, but entail different tasks and so do not directly correspond to productivity.

division of tasks within the team rather than multiple distinct productive abilities of team members.

This paper contributes to a growing literature on the identification of team skills. Arcidiacono et al. (2017), Oettl (2012), and Brave et al. (2017) separately identify a given worker's direct contribution to team output from productivity spillovers onto teammates using worker-level intermediate outcomes of teamwork. The current paper complements these by using productivity measurements of team members taken from outside the team setting, occluding the possibility that individual productivity measurements are influenced by teammates or that they are endogenous to a team level optimization process that assigns different roles to different workers. The results speak also to an emerging literature studying interpersonal skills that exploits qualitative descriptions of job tasks from the Occupational Information Network (O*NET). See Kambourov et al. 2015, Deming 2017, Lise and Postel-Vinay 2016.[9] Interpersonal skills are a primary driver of labour market outcomes, and have been connected to wage polarization (Deming 2017, Autor and Dorn 2013) and the emergence of female labour supply (Cortes et al. 2018, Ngai and Petrongolo 2017). This paper complements this literature by exploiting the direct and clear distinction in job tasks given by singles and doubles pro tennis.

Besides contributing to a growing literature on team skills, the current paper makes a methodological contribution to the literature on two-way fixed effects decomposition of outcomes in connected sets. By comparing value added measurements of a given worker in different settings, I provide a way to identify different skill factors. Benson et al. (2018) is a contemporaneous paper that employs a comparable identification strategy to estimate sales and managerial skill. The authors show that the best sellers are not the best managers, but firms promote successful salespeople perhaps so as to overcome a principle-agent problem. The results of this paper and the current one have implications for research on optimal team size and the division of tasks within an organization.

The remainder of the paper is organized as follows. Section 2 introduces the empirical framework and Appendix A shows identification. Section 3 details the pro tennis dataset and section 4 presents results. Finally section 5 concludes.

[9] Interpersonal skill overlaps conceptually with team skill, but the two are not identical. For example, a gregarious bartender may be said to have good interpersonal skill to the extent that he enthralls customers, but poor team skill if he interacts poorly with other members of staff.

2 Empirical model

2.1 Setting

Consider a set $\mathcal{I}$ consisting of I heterogeneous workers observed over a set of time periods $\mathcal{T}$. Each worker i is characterized by general skill $\alpha_i \in \mathbb{R}$, team skill $\gamma_i \in \mathbb{R}$, and a vector of time-varying observables $X_{it} \in \mathbb{R}^{k_1}$. Summarize a worker i at time t by the triplet $(\alpha_i, \gamma_i, X_{it})$.

Workers produce both individually and in two-person production teams. Solo production depends on a worker's general skill α_i and individual time variables X_i; collaborative skill γ_i plays no role in solo production by definition. Teams consist of workers drawn from a common population.[10] Team output depends on the general skills of either partner α_i, α_j as well as collaborative skills γ_i, γ_j and each partner's time variables X_{it}, X_{jt}. Note that there exist observable properties of teams that are undefined in the solo context. These include team tenure and measures of similarity or difference between partners. Group these team-level observables along with individual-level observables into a of $k_2 \times 1$ vector X_{ijt}, where $k_2 > k_1$. Then a team is characterized by a sextuplet $(\alpha_i, \alpha_j, \gamma_i, \gamma_j, \psi_{ij}, X_{ijt})$, where ψ_{ij} is the idiosyncratic match quality of the team.

2.2 Solo production

Solitary workers produce output according to the following function:

$$y_{it} = \mu_1 + \alpha_i + X_{it}\beta_1 + \varepsilon_{it} \tag{1}$$

where ε is an mean-zero error term. Without loss of generality, make the following normalizations:

$$\sum_{i \in \mathcal{I}} \alpha_i = 0, \qquad \iota_{k_1}\beta_1 = 0$$

where ι_n is a $1 \times n$ vector of ones. Then the constant term μ_1 is the mean of y_{it}. Interpret α_i as the time-invariant component of i's latent productivity. The $I \times 1$ vector $\hat{\alpha}$ of latent productivity parameters can be recovered by performing the classic within transformation and recovering the fixed effects.

[10] This is a non-partite matching setting, wherein any worker i can match with any worker j, $j \neq i$. Both i and j could potentially match with a common partner k. In contrast, typical matching settings are bipartite, consisting of two categories of agents who only match across category, such as firms and workers.

2.3 Assignment

In addition to producing alone, workers match to form teams according to some function $J(i,t)$.[11] This maps $(i,t), i \in \mathcal{I}, t \in \mathcal{T} \to j$, where $j \in \mathcal{I}, j \neq i$, allocating a worker i in period t to a partner from the common population $\mathcal{I}$ in that same time period. The matching function determines set of extant matches, given by $\mathcal{M}$, as well as the duration of each match in the case ongoing partnerships over which $J(i,t) = J(i,t-1)$ for some i,t.[12] An extant match ij is a set of two workers i and j, $i \neq j$, such that $ij \in \mathcal{M}$. Denote the number of extant matches observed as $M \leq \frac{I(I-1)}{2}$.

2.4 Team production

First consider a general model of team production given by

$$y_{ijt} = \mu_2 + \Phi_{ij} + X_{ijt}\beta_2 + \varepsilon_{ijt} \tag{2}$$

where Φ_{ij} is the time-invariant component of match output and is defined by

$$\Phi_{ij} \equiv f(\alpha_i, \alpha_j, \gamma_i, \gamma_j) + \psi_{ij}$$

with ψ_{ij} being the idiosyncratic component of match value. This captures all across-match output variation, and in doing so provides a ceiling for how much a model based on time-invariant parameters α_i, γ_i can explain. Impose the following assumption.

Assumption 1 - *Symmetry*
The team production function y_{ijt} is symmetric in worker inputs; that is, $y_{ijt} = y_{jit}$. Moreover, each component of the team production function is symmetric: $f_{ij} = f_{ji}$, $\psi_{ij} = \psi_{ji}$, $X_{ijt}\beta_2 = X_{jit}\beta_2$.

The interpretation is that there is no specialization of roles; each partner must make qualitatively similar contributions to the team. It will be maintained throughout.[13]

Note that the above functional form places no restriction on the relationship between Φ_{ij} and Φ_{ik} (nor Φ_{jl} for that matter), because the idiosyncratic component of match

[11] Consistency requires that if $J(i,t) = j$ then $J(j,t) = i$. This implies that $J(.,t)$ is an involutory function for all time periods t; that is, $J(.,t) = J^{-1}(.,t)\ \forall t$.

[12] Since the matching is a function of player identity and time, assignment can depend arbitrarily on any characteristic of i. However, it cannot depend on features of the match ij such as idiosyncratic match quality ψ_{ij}. See Assumption 4.

[13] The symmetry assumption is innocuous in the context of the current application to doubles tennis. The rules of tennis require that each partner performs the same range of tasks. See Appendix G for a discussion.

value ψ_{ij} can vary arbitrarily across partner combinations. The lack of restriction eases identification: so long as a combination of partners ij is observed, their match effect Φ_{ij} is estimable. There is no requirement of mobility of workers across teams. Estimation yields an $M \times 1$ vector of parameter estimates $\hat{\Phi}$, which is unbiased under the weak assumption that the error term ε is orthogonal to all regressors.

With the gain of flexibility and weak requirements of the data for identification comes the cost of interpretation: the model given by equation (2) measures whether a team is better than another, but cannot be used to make statements about individual team members. To make such a comparison, introduce the following assumption.

Assumption 2 - *Separability*
The systematic component of the team production function $f(.,.,.,.)$ is separable in worker inputs. It can then be written as $f(\alpha_i, \alpha_j, \gamma_i, \gamma_j) \equiv g(\alpha_i, \gamma_i) + g(\alpha_j, \gamma_j)$.

Since the dependent variable can be measured in levels or logs, this imposes an additive or multiplicative production function respectively. This is equivalent to assumptions made in the wage determination literature (see Abowd et al. 1999, Card et al. 2013) and the teacher value-added literature (see Chetty et al. 2014). It allows a two-way fixed effect decomposition of inputs corresponding to i and j, providing a basis for comparison of individuals rather than only teams. This assumption will be maintained for the baseline results, but it is relaxed in Appendix D.

Denote a worker i's net separable contribution to output as θ_i, where $g(\alpha_i, \gamma_i) = \theta_i$. Then the separable team production function is given by

$$y_{ijt} = \mu_2 + \theta_i + \theta_j + X_{ijt}\beta_2 + r_{ijt} \tag{3}$$

where the residual r_{ijt} now contains the idiosyncratic component of match productivity ψ_{ij} in addition to the error term ε_{ijt}; that is,

$$r_{ijt} = \psi_{ij} + \varepsilon_{ijt}$$

which requires that the idiosyncratic match effect ψ be uncorrelated with observables in order for unbiasedness of parameter estimates. Intuitively: if $\theta_i, \theta_j \perp \psi_{ij} | ij \in \mathcal{M}$, then idiosyncratic match quality is unrelated to the assignment process. In this case θ_i gives i's value added to a randomly assigned team, and can be interpreted as an input in the production function (label this 'composite skills' – general and team skills aggregated into a single skill index). However, if workers anticipate match quality and choose partners conditional on high quality, θ_i should be interpreted as i's input into production plus her ability to seek out good matches; placing i in a random team will not increase that team's output by θ_i. See Appendix A for a full discussion and subsection 4.8 for evidence on the exogeneity of ψ to the assignment process.

As with solo production, impose the normalization assumptions

$$\sum_{i \in \mathcal{I}} \theta_i = 0, \qquad \iota_{k_2} \beta_2 = 0$$

which imply that μ_2 is the mean of y_{ijt} across dimensions. Note also that assumption 2 and equation (3) imply the following relationship between match effects and composite skills:

$$\Phi_{ij} = \theta_i + \theta_j + \psi_{ij} \tag{4}$$

which decomposes the match effect into individual inputs and the idiosyncratic match effect.

Equation (3) provides a familiar framework for comparing workers' contributions to output in a collaborative setting, having been used to analyze teacher value-added and worker and firm wage determinants. Such settings as classrooms and the broader labour market typically present an outcome (student test scores, wages) that is the product of multiple inputs (students and teachers, workers and firms) and uses functional forms similar to equation (3) to decompose the contribution of the two groups of inputs.[14]

The current paper innovates by considering solo production as given by equation (1) in conjunction with team production. This introduces a novel dimension of variation: within-task, within-individual, across-team size. This measures to an individual's relative capability at solo versus team production. Impose the following functional form restriction on a worker's net contribution to output:

$$\theta_i = a\alpha_i + \gamma_i \tag{5}$$

where parameter a determines the transferability of general skills into team production. Equation (5) implies the following definition of team skill.

Definition 1 *For a given worker i, define team skill γ_i as the component of composite skill θ_i that is not explained by general skill α_i.*

[14] Typically the two-way fixed effects decomposition has been applied to the bipartite matching settings noted above, in contrast to the current non-partite setting. For a discussion of relevant concerns see sections A and C.

Then γ_i is orthogonal to α_i by construction: it is a worker i's comparative advantage in team production compared to solo. The functional form given by equation (5) will be generalized in Appendix D.

Combining the team production function given by equation (3) and the functional form of composite input $g(.,.)$ given by equation (5) yields team production function

$$y_{ijt} = \mu_2 + a(\alpha_i + \alpha_j) + \gamma_i + \gamma_j + X_{ijt}\beta_2 + r_{ijt} \tag{6}$$

which illustrates the role of own and team skill in team production. Equation (6) has equal explanatory power to (3); they are identical except that (6) decomposes worker inputs into general and team skill components.

Note that in team output data, α_i and γ_i do not vary disjointly – they are both present in every team i participates in. Separate identification of these parameters requires observation of solo production. Estimation can proceed in one of two ways, yielding identical parameter estimates. General skills α can be recovered from fixed-effects estimation of equation (1), and the parameter estimates included as data in the estimation of team production function (6).[15] Alternatively, estimation of equations (1) and (6) can be done by seemingly unrelated regression (or separately), yielding estimates of general skills α and composite skills θ, with team skills γ being recovered in a second step by estimation of auxiliary regression equation (5). The point estimates of parameters produced by either method are identical. See Appendix C for details.

2.5 Variance decomposition

The main results of the current paper will be the share of team output explained by team skill γ.[16] Because equation (2) includes a fixed effect for each extant ij combination, it captures all across-team variation. Note that equations (4) and (5) imply that match effect can be decomposed as follows:

$$\Phi_{ij} \equiv a(\alpha_i + \alpha_j) + (\gamma_i + \gamma_j) + \psi_{ij} \tag{7}$$

which divides match effects into three terms: the first given by the general skills of both partners, the next given by team skills of both partners, and the last by the idiosyncratic component of team productivity.

[15] This may introduce measurement error. See Appendix C for a discussion, and subsection 4.4 for evidence that measurement error is not severe in the current application.

[16] Variance share can be defined as a share of total output, or as a share of across-team output variation. The current subsection covers the latter; for a discussion of the former see Appendix E.

Equation (7) implies that across-match variation can be decomposed as follows:

$$\sigma_\Phi^2 = a^2\sigma_\alpha^2 + 2a\sigma_{\alpha\gamma} + \sigma_\gamma^2 + \sigma_\psi^2 \tag{8}$$

with terms corresponding to own and team skill regressor groups (for both partners i and j), the covariance between the two regressor groups, and the idiosyncratic match effect.[17] According to the above notation $\sigma_\alpha^2 = \mathrm{var}(\alpha_i + \alpha_j)$ and $\sigma_\gamma^2 = \mathrm{var}(\gamma_i + \gamma_j)$. In other words,

If the greater portion of across-match variation is explained by general skills as measured by solo performance, the $a^2\sigma_\alpha^2$ term should be large. On the other hand, a large estimate of σ_γ^2 says that a worker-fixed factor unexplained by solo performance – defined as team skill – is the major determinant of a team's productivity. Finally, if σ_ψ^2 is large then the separable production function is a poor model of the data generating process – that a team's performance cannot be well approximated by the performance of its members on other teams.

The variances of regressor groups can be further decomposed. First consider the variance component attributed to general skills, given by σ_α^2. It decomposes into

$$\sigma_\alpha^2 = \sigma_{\alpha_i}^2 + 2\sigma_{\alpha_i\alpha_j} + \sigma_{\alpha_j}^2 \tag{9}$$

where $\sigma_{\alpha_i}^2$ denotes the variance in general skills of the arbitrarily chosen reference worker, and $\sigma_{\alpha_j}^2$ of the equally arbitrarily identified partner. In a bipartite matching setting (such as workers and firms) these would correspond to qualitatively different objects, and may differ in magnitude; in the current setting they represent skill factors attributed to groups of workers from the same population $\mathcal{I}$, and should be expected to be equal if reference identities i and partner identities j are chosen randomly and the sample size is large. The remaining covariance component indicates whether high general skill workers match with others of their like (positive assortative matching), with low general skill workers (negative assortative matching), or if there is no discernible sorting pattern.

Similarly, decompose team skill's variance component as follows.

$$\sigma_\gamma^2 = \sigma_{\gamma_i}^2 + 2\sigma_{\gamma_i\gamma_j} + \sigma_{\gamma_j}^2 \tag{10}$$

Analysis of the components follows that of σ_α^2 above.

[17] Recall that the latter is assumed to be uncorrelated with observables, and so introduces no extra covariance terms.

3 Data

This section describes the professional tennis dataset. It takes the form of an unbalanced panel of tournament results from men's singles and men's doubles events.[18] It resembles a typical personnel economics dataset in that it details high frequency productivity data on individual workers. The doubles observations resemble a matched employee-employer dataset in that the identities of either partner are reported and can be matched across observations. However, in this case either party to a match is drawn from a single population (non-partite matching), in contrast to the typical case wherein agents from one group match to agents from a mutually exclusive group (bipartite matching).[19]

The following subsections describe the labour market for professional tennis and then the structure of the data. The next reports summary statistics and gives a portrait of the data, highlighting the resemblance of lifecycle earnings and partnership dynamics in professional tennis to those of the broader labour market. For more information on the rules of tennis, see Appendix G, and for supplementary information recovered from interviews with tennis professionals see Appendix H.

3.1 Organizational structure of professional tennis

The Association of Tennis Professionals (ATP) organizes the ATP World Tour, the world's elite men's professional tennis circuit. The ATP maintains the worldwide rankings that dictate eligibility to their own tournaments as well as to the annual cups and grand slams organized by the International Tennis Federation (ITF). Grand slams consist of the Australian Open, Roland Garros, Wimbledon, and the U.S. Open; these award 2000 ranking points to champions, and a declining series of points awarded to runners up and lower achieving competitors.[20] ATP tournaments range from the mandatory Masters 1000 series tournaments – of which there are nine per year, winners being awarding an eponymous number of ranking points – to the 500 and 250 series. Rankings are determined by the number of points accumulated in tournaments played during the last 52 calendar weeks.

In addition to the aforementioned ATP World tour (consisting of the 1000, 500,

[18] The professional women's circuit is run by a different organization under different rules (although some events are joint). For this reason the current paper focuses on men's tennis.

[19] An example of a bipartite matching setting is the typical labour market setting wherein workers match with firms and no 'within-group' matching is possible.

[20] Runners-up typically earn 50% of the ranking points of winners, with points progressively declining by about half with each step down.

and 250 point tournaments) and the four grand slams, there are three other classes of tournament. The ATP maintains the ATP Challenger Tour, which is a level below the ATP World Tour in that the ranking eligibility requirements are less stringent. Players are ranked on the same scale as on the ATP World Tour, so Challenger tournament winners earn ranking points that determine eligibility also to World Tour events. Below the Challenger Tour is the ITF Men's circuit, which holds Futures tournaments, whose points are also incorporated into the ATP rankings.

All tournaments mentioned above consist of both singles and doubles play categories. Most players participate in both categories; this paper restricts its attention to these players for the main analysis. The ATP maintains separate rankings for singles and doubles play. While eligibility for singles tournaments is based simply on one's currently held number of ranking points, eligibility for doubles tournaments is determined by summing the ranking points held by the two team members.[21] Prize money is awarded equally to each player on a team, except in the case of withdrawal, in which case the withdrawing player is penalized by the tournament organizers. Side payments are unheard of, and would likely be considered unsportsmanlike.[22]

3.2 Data structure

The ATP dataset includes tournaments sanctioned both by itself and by the ITF. Together these comprise the entire professional history of men's tennis players. The singles data resemble a labour force survey in that they report an unbalanced panel of earnings by tournament. For doubles play, teammates are matched, resembling a matched employee-employer dataset. Within a tournament, teams do not vary in terms of membership. Many doubles teams last one season or more, but most do not last this long. Players are frequently mobile across doubles partners; the largest connected set consists of over 93% of players and 97% of tournaments played (with the more connected players tending to play more tournaments). For the main results, the sample is limited to players who participate in at least 20 singles and 20 doubles tournaments during the sample period. In this case, the largest connected set contains all observations.

The basic unit of observation is a player-tournament in singles and a team-tournament in doubles. An observation consists of one or two player identities, observable player and team characteristics, the date the tournament took place, and an outcome. Player

[21] 2017 ATP Official Rule book, VII 7.13 C.

[22] See Appendix H for details.

Table 1: Summary Statistics

	All tournaments					
	Singles prize	Doubles prize	Age	Experience	Year	Month
mean	4648	780	24	65	2012	6
p50	292	141	23	54	2012	6
sd	42838	4238	4	48	2	3
min	0	0	14	1	2009	1
max	3800000	285000	42	249	2015	12
count	172779	50603	273985	273985	223382	223382

Prize money is reported in 2015 USD, per player (each member of a doubles team receives the amount listed). The number of singles observations is $N_1 = 172779$, and the number of doubles observations is $N_2 = 50603$. Age and experience then are reported for $N_1 + 2 \times N_2 = 273985$ observations – one per singles player-tournament, and two per doubles team-tournament. Year and month are reported for $N_1 + N_2 = 223382$ observations.

characteristics include age, tournament experience, height, weight, backhand style, and handedness. The observable team characteristic is team tenure (number of tournaments; interactions of player characteristics are also considered).

The preferred measure of outcome is prize money. Alternatives include ranking points and round achieved (winner, runner up, etc.). Ranking points track prize money closely, and using them would avoid currency conversion and inflation issues. The set of possible rounds achieved varies for tournaments of different sizes – a winner has reached the fifth round of a 32-person tournament, but the seventh round of a 128-person tournament. Prize money is preferable because it has a tangible economic implication both to the reader of this paper and to the player who relies on it to continue playing tennis professionally.

Prize money varies by tournament and across years, moving roughly in proportion to inflation. Payments are made in USD, GBP, AUD, or EUR. This paper considers real earnings in 2015 USD, using CPI and exchange rate data from Penn World Tables. Reliable prize money data is available from 1985 onward. A series of reorganizations resulted in a payment scheme that is consistent from 2009 to present (see Gibson 2010). As this caused a change in the distribution of prize money, the main results of this paper are limited to the sample ranging from 2009 to 2015.

3.3 Summary statistics

Table (1) summarizes the key variables from the ATP dataset. One observation corresponds either to a player-tournament in the case of singles, or a team-tournament in the case of doubles. As mentioned above, the sample is restricted to tournaments occurring from January 2009 to December 2015, so as to guarantee consistency of payment schedules. In addition, the baseline sample omits players observed for fewer than 20 tournaments in either category so as to prevent bias from measurement error (see Appendix C for a discussion).

Note that although the mean prize money won per tournament is greater in the singles category, the median is of the same order – the median singles player wins $297 per tournament while the median team wins $282 (with $141 awarded to each player). Tennis being a physical profession, players enter the industry at an earlier age and retire earlier.

Figure 1 shows the distribution of yearly earnings. The right-skewness is in part an artifact of the prize money schedule, but is also endogenous to players' participation decisions and tournament outcomes.

Several emergent properties of the ATP dataset mirror their counterparts in the broader labour market. This speaks to the fact that tennis is a profession rather than just a recreational activity, and facilitates comparison of professional players with a 'representative worker' from a typical economy. The lifetime average earnings profile shown in figure 2 exhibits the familiar inverted U shape of the typical worker in the broader labour market. Due to the athletic nature of the profession, earnings peak earlier – around age 30 – and then dip before retirement. This feature is slightly attenuated in doubles competition, with earnings peaking and workers retiring slightly later than in the singles category. Doubles earnings are consistently below singles – following the prize money schedule – until near retirement, when they overtake singles earnings. Figure 3 shows the distribution of the first-difference of yearly earnings in both play categories. Interestingly, year-to-year earnings shocks follow a similar non-normal distribution with excess kurtosis as that in the broader US economy as described in Guvenen et al. (2015). The singles and doubles distributions look nearly identical, indicating a comparable level of earnings uncertainty across categories.

In addition to the career of a tennis player looking broadly similar to the career of a typical worker, doubles teams resemble job matches in the broader labour market. Figure 4 shows the hazard rate of separation declining as partners play together longer,

Figure 1: Distribution of yearly log earings

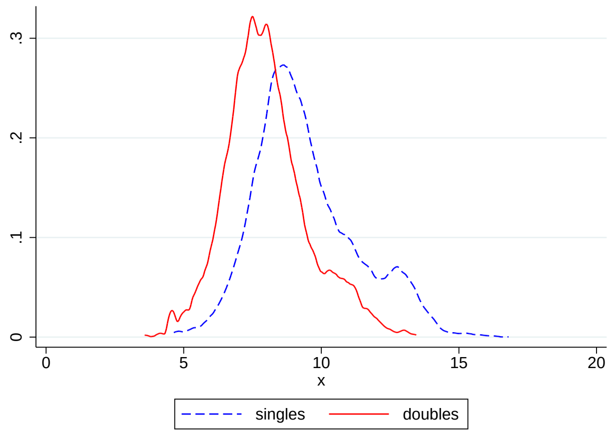

Figure 2: Average log earnings by age

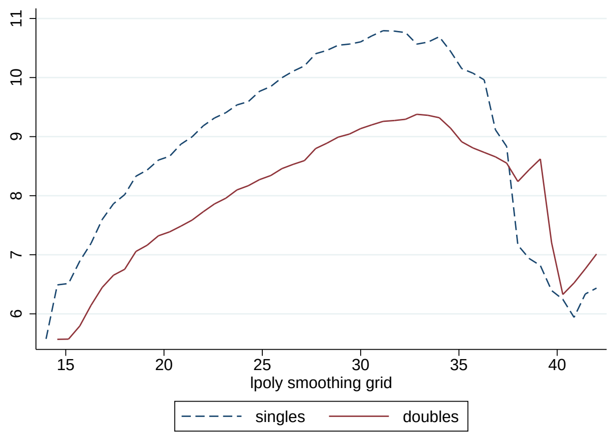

Figure 3: Year to year log earnings changes

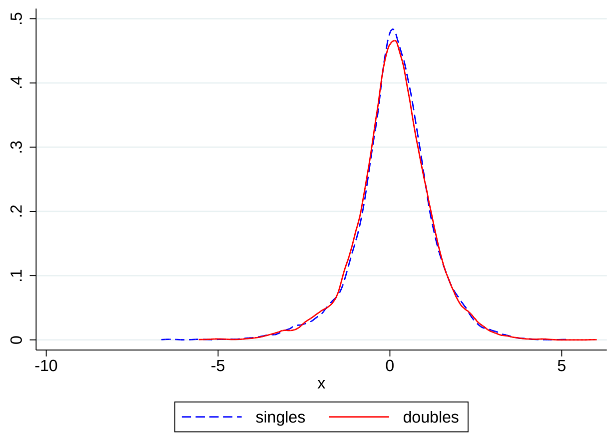

Figure 4: Separation hazard rate by team tenure

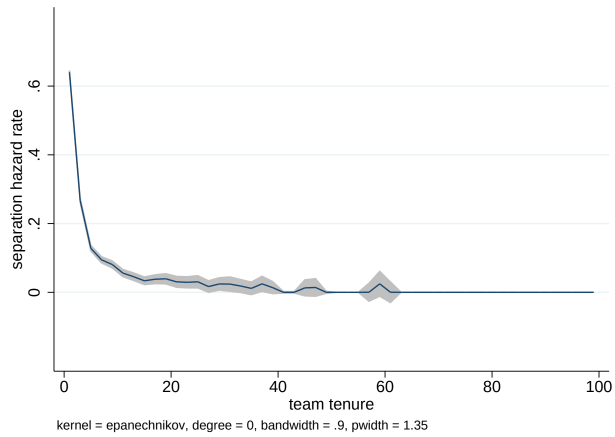

Figure 5: Average log earnings by team tenure

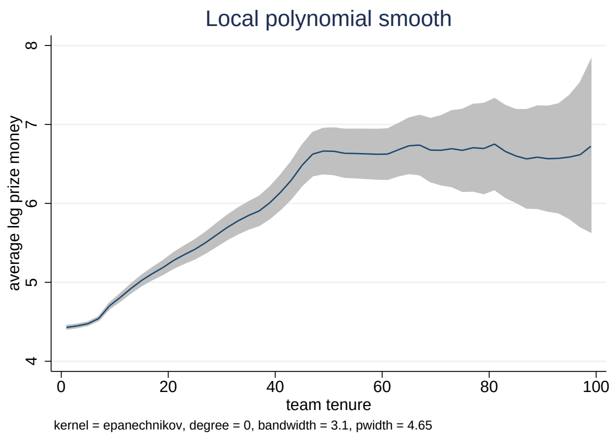

and figure 5 shows an increasing, concave relationship between average log earnings and team tenure, as in the broader labour market (see for example Buhai et al. 2014).

4 Results

The primary result is that there is that there is a measurable second factor active in team production; that is, the identity of a team member is a large and statistically significant predictor of team output, even when controlling for general skill of either member (as estimated from solo production). This is portable across teams, and unexplained by observables and performance in solitary production. The following subsection details the statistical and economic significance of this second factor. This finding is robust to a flexible nonlinear functional form specification of how general skill contributes to team production (see the nonlinear model given by equation 26). Moreover, the interactive term in the nonlinear model is significant statistically but not economically, and does not contribute to the explanatory power to the model (neither R^2 nor $\bar{R}^2$ improves compared to the linear model). This validates the log-separability assumption of the widely applied two-way fixed effects decomposition in the current setting. Furthermore, an unrestricted match-effects model does not significantly outperform the two-way additive model when controlling for number of parameters used. A lack of systematic sorting patterns related to the match effect residual validates the assumption of the two-way fixed effects model that assignment is orthogonal to the unobserved idiosyncratic component of match value. See subsection 4.8 for more on this.

Subsequent subsections illustrate matching patterns between partners: individuals sort according to both general skill and team skill, but there is no systematic sorting pattern across skill dimensions.

4.1 Model fit

Table 2 summarizes the four regression models: solo production, team production, nonlinear team production, and the saturated match effects model. Specifications with and without observable controls are reported; most of the explanatory power comes from the fixed effects. The team and nonlinear team production and saturated models are estimated on the doubles sample. All models explain a significant share of output variation. Introducing nonlinear terms to team production doesn't add significantly to explanatory power. In absolute terms the saturated match effects model fits best, but controlling for

Table 2: Model fit

	Singles	Singles	Doubles	Doubles	Nonlinear	Saturated	Saturated
α	1	1	0.358***	0.338***	0.233***	-	-
α^2	-	-	-	-	0.0756	-	-
$\alpha_i \times \alpha_j$	-	-	-	-	-0.00447*	-	-
Ctrls.	-	x	-	x	x	-	x
Obs.	170361	170361	50035	50035	50035	50035	50035
R^2	0.38	0.39	0.4	0.46	0.46	0.68	0.7
$\bar{R}^2$	0.38	0.39	0.38	0.43	0.43	0.4	0.42

The singles specification is given by equation (1) and doubles by equation (6). The saturated model is given by equation (2). The nonlinear specification is given by equation (29), found in Appendix D.

the number of parameters, it does no better than the log-separable specification (see the adjusted r-squared statistic $\bar{R}^2$). A joint significance test of the match effects fails to reject the two-way model. It appears the better fit stems from flexibility; the saturated model has around 20000 fixed effects parameters while the baseline model has around 2000.

The general skill parameter vector α is measured in terms of output in log USD. By construction one unit of general skill returns one unit of output in solo production. When included as a regressor in team production, the general skill fixed effect returns less than a unit of output in team production, giving some sense of how transferable are skills between solo and team production (see table 2). Note that the constant term included in each regression specification makes this transferability measure scale invariant: if all doubles prized were doubled, the coefficient estimate would be identical. The nonlinear specification of team production suggests that the returns to general skill in team production may be convex, and that general skill inputs are log-submodular across partners. However, these effects are small in magnitude, and do not significantly improve explanatory power.

4.2 Variance decomposition of match effects

The last two columns of table 2 show that the saturated match effects model explains around 70% of total output variation for doubles matches. Match effects capture all

Table 3: Variance decomposition of match effects

total		general skills		covar.		team skills		resid.
σ^2_Φ	=	$a^2\sigma^2_\alpha$	+	$2a \times \sigma_{\alpha\gamma}$	+	σ^2_γ	+	σ^2_ψ
1	=	0.2	+	-0.07	+	0.5	+	0.38
omitting control variables:								
1	=	0.21	+	-0.09	+	0.48	+	0.4

This table decomposes the variance of the match effect (given by $\Phi_{ij} \equiv a(\alpha_i + \alpha_j) + \gamma_i + \gamma_j + \psi_{ij})$ into its components according to the definition of the variance. The residual match effect ψ_{ij}, being assumed orthogonal, does not introduce additional covariance terms. See subsection 2.5 for details.

across-match variation; these models provide a ceiling for how much variation the two-way models can explain. Table 3 normalizes across-match variation to one, and decomposes it into components explained by variation in general skilll α, team skill γ, covariance between the two, and the remaining residual match effect (note that the residual is orthogonal to observables so there are no additional covariance terms). Variation in team skills accounts for 50% of across-team output variation. A smaller portion of 20% of variation is attributable to general skills, and a good portion is attributed to idiosyncratic variation at the team level. The covariance between the general skill and team skill regressor groups is small and negative.[23]

Now, further decompose the variance shares of regressor groups into components attributable to either partner, and to the share given by the covariance of partner skills. Tables 4 and 5 decompose general and team skill regressor groups respectively. Note that i and j identities are assigned to partners alphabetically; since there is no reason to expect a difference between the variance of their fixed effects, in large samples their variances should be equal. This holds in the current sample. For both skill classes, roughly two-thirds of the variance is attributed to variation in the skill levels of either partners, with one-third attributed to the positive covariance between partners' skills.

4.3 Variance decomposition of all regressors

This subsection decomposes total output variation, rather than only across-match variation. The measurements of variance decomposition are in line with those used by the

[23]While own and team skill *within and individual* are orthogonal by construction, this is not the case *across partners*. Thus $\sigma^2_{\alpha\gamma} = \text{cov}(\alpha_i, \gamma_j) + \text{cov}(\alpha_j, \gamma_i)$.

Table 4: Variance decomposition of general skill effects

total		partner i		covar.		partner j
σ^2_α	$=$	$\sigma^2_{\alpha_i}$	$+$	$2 \times \sigma_{\alpha_i \alpha_j}$	$+$	$\sigma^2_{\alpha_j}$
1	$=$	0.34	$+$	0.32	$+$	0.34
			without controls:			
1	$=$	0.38	$+$	0.35	$+$	0.38

This table decomposes the variance of general skills α_i, α_j within a team ij according to the definition of the variance.

Table 5: Variance decomposition of team skill effects

total		partner i		covar.		partner j
σ^2_γ	$=$	$\sigma^2_{\gamma_i}$	$+$	$2 \times \sigma_{\gamma_i \gamma_j}$	$+$	$\sigma^2_{\gamma_j}$
1	$=$	0.33	$+$	0.34	$+$	0.33
			without controls:			
1	$=$	0.36	$+$	0.33	$+$	0.36

This table decomposes the variance of team skills γ_i, γ_j within a team ij according to the definition of the variance.

literature on two-way fixed effects decomposition of wage determinants.[24]

Table 6 presents three measurements of variance decomposition based on a singles regression estimated by equation (1). These are the balanced variance share (BVS),[25] the simple correlation between a group of predicted regressors and output, and a predicted regressor's standard deviation. The latter two measures are preferred by AKM. See Appendix E for a comparison with other variance decomposition measures.

According to all methods of variance decomposition the most important predictor of singles outcomes are player fixed effects α. These explain 34% of output variation. Time-variables experience (number of tournaments played), age (in years), and calendar year and month are also significant predictors. These results are invariant to the inclusion of tournament fixed effects (see Appendix E).

Table (7) presents a variance decomposition of doubles tournament results based on the linear specification (6). The general skill fixed effect regressors are those predicted by equation (1) and explain 12% of output variation. Crucially, team skill fixed effects γ are the most important predictor of outcomes for all variance decomposition measures, explaining over a third of output variation in the linear specification and over one-quarter in the nonlinear. Individuals posses a fixed factor that adds value to every team they participate in, that is not explained by their performance in singles. In both linear and nonlinear specifications (Tables 7 and 8) these fixed effects are highly statistically significant, as evidenced by the F-statistics and corresponding p-values of less than 0.01.

Note that the general skill terms (associated with α) have an F-statistic of zero and so appear insignificant by this measure. This is because the F-statistic is generated by comparing a restricted model excluding the $\hat{\alpha}_i$ and $\hat{\alpha}_j$ regressors to the unrestricted model; since both contain fixed effects corresponding to the identities of i and j, the fixed effects in the restricted model absorb all explanatory power given by $\hat{\alpha}$. The positive BVS measurement, as well as the regressor group's correlation with outcomes and standard deviation, indicate its economic significance in determining team output;

[24] See for example Abowd et al. (1999), Card et al. (2013).

[25] The BVS is defined as follows. Consider a regression model $y = x_1\beta_1 + x_2\beta_2 + u$, with x_i being an $n \times k_i$ matrix corresponding to a grouping of $k_i \geq 1$ regressors for $i = 1, 2$. Assuming orthogonality of the error term, this implies $\text{var}(y) = \text{var}(x_1\hat{\beta}_1) + 2\text{cov}(x_1\hat{\beta}_1, x_2\hat{\beta}_2) + \text{var}(x_2\hat{\beta}_2) + \text{var}(\hat{u})$. The BVS of a regressor x_1 is defined as

$$\text{BVS}(x_1) \equiv \frac{\text{var}(x_1\hat{\beta}_1) + \text{cov}(x_1\hat{\beta}_1, x_2\hat{\beta}_2)}{\text{var}(y)} \tag{11}$$

which apportions covariance between regressors equally across regressor groups. The BVS of all regressor groups sums to the total R^2 of the regression.

Table 6: Variance decomposition for singles regression

	BVS	$\rho(., Y)$	Std(.)	F	p
α	0	0	0	0	1.0
experience	0.34	0.6	1.4	41.91	0
age	0.02	0.21	0.25	0	1.0
year	0.02	0.21	0.25	0	1.0
month	0	0.02	0.36	50.05	0
tourney	0	0.06	0.18	131.13	0

This table reports alternative measurements of variance decomposition for general skills α and time-varying observables for the singles specification given by equation (1). These are the balanced variance share (BVS), the raw correlation between a predicted regressor group and output (ρ), and the standard deviation of the predicted regressor group. See Appendix E for details. Also reported are the F-statistics for a joint significance test of the regressor group and the corresponding p-values.

Table 7: Variance decomposition for doubles regression – linear

	BVS	$\rho(., Y)$	Std(.)	F	p
α	0	0	0	0	1.0
γ	0.1	0.29	0.81	0	1.0
experience	0.28	0.51	1.26	7.47	0
tenure	0.03	0.21	0.32	1.36	0
age	0.01	0.13	0.11	0.81	0.97
year	0	0.06	0.19	1.14	0.07
month	0.03	0.16	0.45	190.22	0
tourney	0	0.03	0.15	24.76	0

Table 8: Variance decomposition for doubles regression – nonlinear

	BVS	$\rho(., Y)$	Std(.)	F	p
α	0	0	0	0	1.0
α^2	0.07	0.29	0.56	0	1.0
$\alpha_i \times \alpha_j$	0.1	0.39	0.6	0	1.0
γ	0	-0.24	0.02	1.26	0.26
experience	0.21	0.49	0.99	5.39	0
tenure	0.03	0.21	0.32	1.36	0
age	0.01	0.13	0.11	0.81	0.97
year	0	0.06	0.19	1.14	0.07
month	0.03	0.16	0.45	189.89	0
cons.	0	0.03	0.15	24.7	0

Figure 6: Baseline estimates of γ versus instrumented estimates

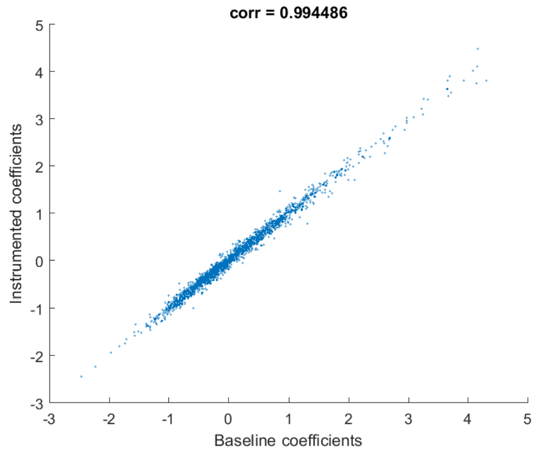

as reported in table 2, the coefficient on α in the linear model is highly significant.

4.4 Testing for measurement error

Using predicted coefficients from singles play may introduce measurement error, as discussed in Appendix C. This subsection describes a simple test to determine the magnitude of measurement error. Measurement error is found to be negligible when the sample is restricted to players observed for at least 20 tournaments in each category over the sample period.

The test is as follows. First split the singles sample into two halves. Within an individual, the first half of observations are allocated to sample 1, and the second to sample 2. Estimate α separately for the two samples according to equation (1). This produces two sets of parameter estimates, $\hat{\alpha}_1$ and $\hat{\alpha}_2$. Now estimate equation (6), instrumenting the $\hat{\alpha}_1$ parameters with $\hat{\alpha}_2$. This produces instrumented estimates of team skill $\hat{\gamma}_{iv}$. These instrumented estimates are plotted against the baseline in figure 6.

The estimates of γ using the split-sample instrument approach are virtually identical to the baseline estimates. This suggests little scope for measurement error in the estimated parameters $\hat{\alpha}$.

4.5 Testing for variation in effort

One concern is that players with a comparative advantage in doubles – those with higher team skill factors – exert more effort in doubles matches compared to singles, and vice

Figure 7: Skill factor estimated using full sample versus only busy tournaments

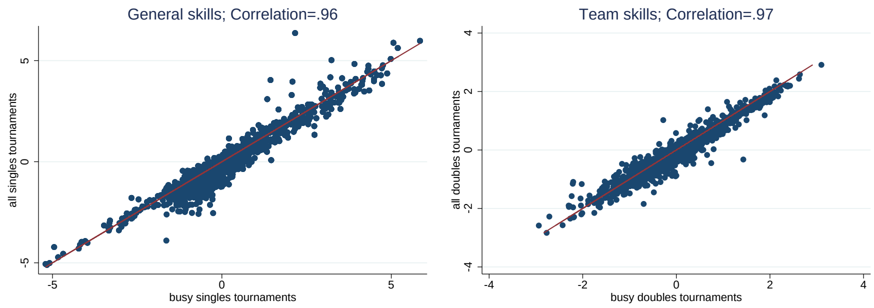

versa. In this case, differences in estimated skill factors are exaggerated by the effect of effort. This subsection presents a simple test for the effects of effort.

In any given tournament, a player may participate in one of the doubles or singles categories, or both. Call a tournament in which a player participates in both categories a 'busy tournament'. If players ration their efforts, one would expect to see a decline in performance during busy tournaments. On the other hand, if players exert full effort in all matches, skill estimates should not differ when all tournaments are considered versus busy tournaments only. Figure 7 reports correlations of baseline estimates of skill factors versus estimates constructed only from busy tournaments.

The estimates are constructed as follows. Baseline estimates using the full sample are plotted on the vertical axes. The 'busy' general skill estimates are recovered from a restricted sample that omits all tournaments in which players only participate in the singles category. If players ration their efforts when they have to concentrate also on doubles play, the busy estimates should be lower. However, the two sets of estimates are nearly identical. Busy team skill estimates are recovered by restricting the sample to busy tournaments for both stages of the regression; so general skills as well as team skills are recovered from busy tournaments only. Again, these are nearly identical to the baseline estimates.

Observed performance does not substantially differ when players are able to focus only on one category of play, suggesting that the rationing of effort does not play a large role in determining output or the skill factor estimates.

Figure 8: Joint distribution of general skills α_i, α_j by team, all tournaments

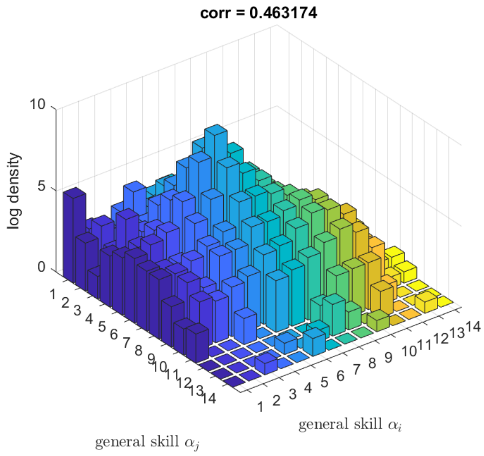

Figure 9: Joint distribution of team skills γ_i, γ_j by team, all tournaments

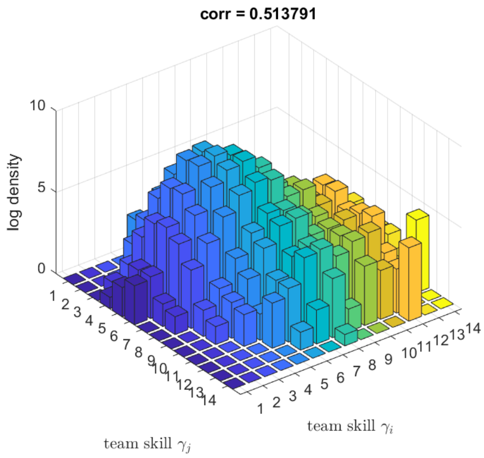

Figure 10: Cross distribution of skills α_i, γ_j by team, all tournaments

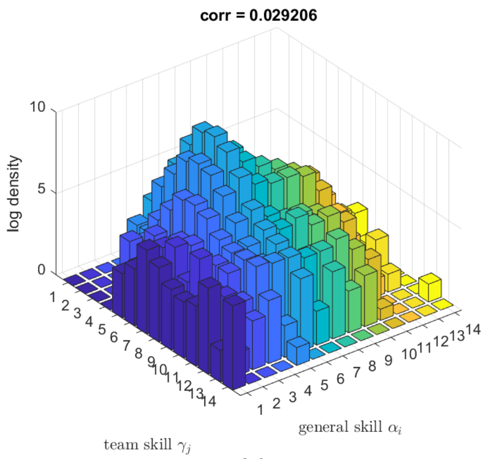

Figure 11: Joint and cross distributions of skills, ATP World Tour + ITF Grand Slams

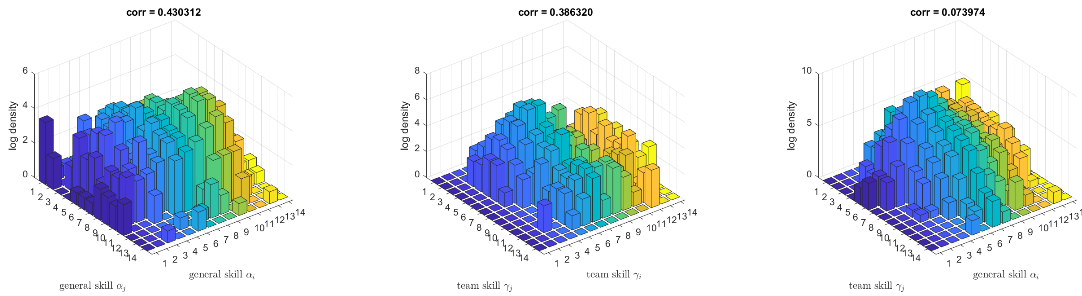

Figure 12: Joint and cross distributions of skills, ATP Challenger Series

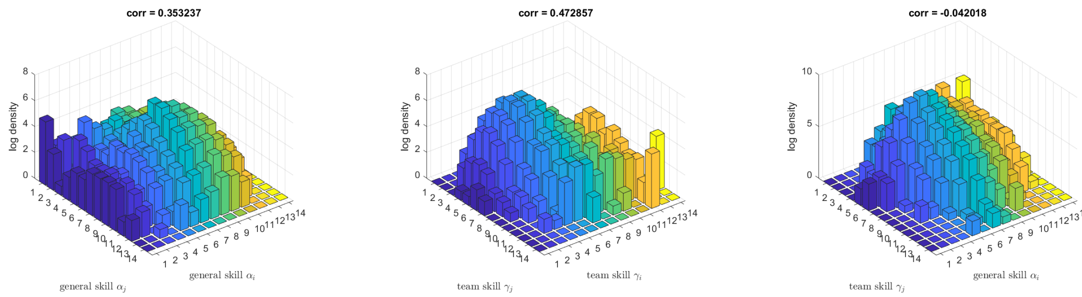

Figure 13: Joint and cross distributions of skills, ITF Futures events

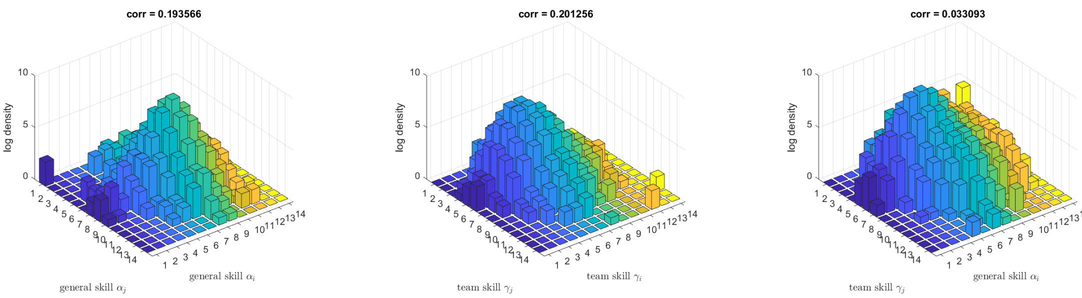

4.6 Sorting patterns

Figure 8 plots the joint distribution of partners' general skills $(\hat{\alpha_i}, \hat{\alpha_j})$ for all extant matches ij. Note that $\hat{\alpha}$ is estimated from singles data, while the set of extant matches comes from doubles; there is no possibility of a mechanical relationship between $\hat{\alpha_i}, \hat{\alpha_j}$. There is a positive association between teammates' general skill levels and a correlation of almost one-half. Figure 17 shows a similar pattern of sorting by team skills $\hat{\gamma}$; although team skill γ_i is the residual player-level fixed effect after the partial effect of general skill α_i has been removed, the correlation between partners' team skills is higher than that of general skills. There is little indication that partners sort across skill classes, as evidenced by figure 10.

Figures 11, 12, and 13 present the same joint distributions for the separate leagues: the world tour and slams, challenger series, and futures series respectively. Within each league the same sorting patterns persist to different degrees. Sorting along the general skill dimension is strongest for the world tour and grows weaker as the level of tournament falls. Sorting by team skills is strongest in the challenger series and weakest in futures events. None of the leagues exhibit strong sorting patterns across skill types.

All sorting results are robust to using the nonlinear specification.

4.7 Returns to skills

Positive assortative matching means that better players are matched with better partners. This implies that in addition to the direct return to skills (a player i with a higher general skill α_i earns more because he is more productive) there is an indirect allocative return to skills (i earns more also by virtue of being matched to a partner who is more productive). Table 9 reports the percentage earnings returns associated with a one standard deviation increase in general skill α and team skill γ, both directly and mediated through one's partner.

General skills provide a large return in the singles category, in the order of 130% higher earnings for a one standard deviation increase.[26] General skills provide a smaller return in doubles at around 38% higher earnings. A player one standard deviation higher in general skill α can be expected to match with a partner who also has a higher general skill, resulting in 20% higher earnings through the allocative channel. Doubles tennis is more intensive in team skill than in general skill: the direct and allocative returns to team skill γ are 66% and 35% respectively.

[26] This is invariant to level shifts in the singles earnings distribution, which are captured by $\hat{\mu_1}$.

Table 9: Direct and allocative returns to skills

	Earnings returns to skills (% per one standard deviation)	
	Singles	Doubles
α direct	124.1	42
α allocative	-	20.7
γ direct	-	66.1
γ allocative	-	34

Direct returns to skills are equal to the earnings return associated with a one standard deviation increase in general skill α or team skill γ, holding the partner constant. Beyond the direct return, players receive an allocative return to skills by virtue of being matched to a better partner.

4.8 Exogeneity of residual match effects

The identifying assumption of the two-way fixed effects decomposition is that the idiosyncratic component of match effects is orthogonal to the regressors. Crucially, the regressors include a dummy variable for partner identity; this means that the matching function $J(i,t)$ does not depend on idiosyncratic match effect ψ.

To test this assumption, consider the residual match effect $\hat{\psi}$. This can be calculated by taking the model residuals $\hat{r}_{ijt}$ and averaging over time within the match. There are residual match effects: some teams consistently outperform what the two-way fixed effects model predicts, and others underperform. However, so long as these match effect residuals do not determine the assignment of team members, the model assumption is valid.

Figure 14 plots the histogram of residual match effect $\hat{\psi}$. The distribution is unimodal and approximately symmetric, with excess kurtosis compared to a normal distribution. Figure 15 plots the realized length of a match – how many tournaments a team played together – against the residual match effect. Two types of sorting pattern would violate the exclusion restriction. First, an increasing association would indicate that better matches remain together for longer; no such pattern is apparent. Second, a truncated left tail would indicate that potential partners evaluate their match productivity in advance, and decline to form team with poor idiosyncratic match quality. There appears to be no evidence of this either. Figure 16 shows average match effect residual by tenure achieved. An increasing association would indicate that teams who stay together longer

Figure 14: Distribution of residual match effect ψ

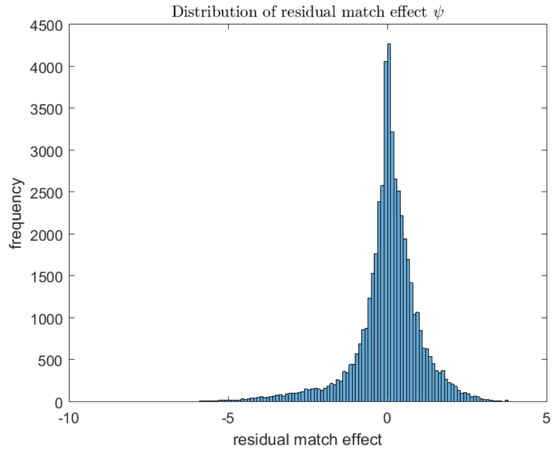

Figure 15: Joint distribution of residual match effect ψ, match tenure

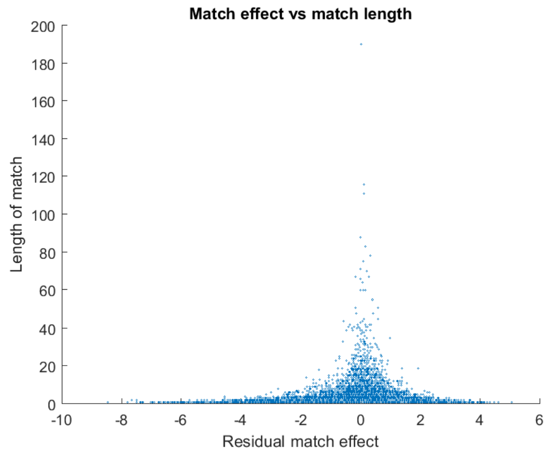

Figure 16: Average residual match effect ψ by achieved match tenure

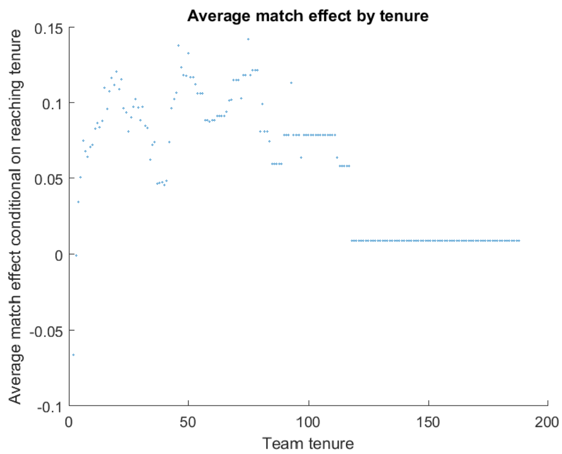

Table 10: Cox proportional hazard analysis of team separation

	coef	exp(coef)	p
$(\alpha_i + \alpha_j) \times a$	0.11	1.12	0
$\gamma_i + \gamma_j$	-0.19	0.83	0
ψ_{ij}	-0.06	0.94	0

This table reports results of a Cox proportional hazards regression on team duration, including as regressors the combined general skill of the teammates $(\alpha_i + \alpha_j) \times a$, combined team skill $\gamma_i + \gamma_j$, and residual match effect ψ_{ij}.

Table 11: Cox proportional hazard analysis of team separation with skill disparities

	coef	exp(coef)	p
$(\alpha_i + \alpha_j) \times a$	0.08	1.09	0
$\gamma_i + \gamma_j$	-0.26	0.77	0
ψ_{ij}	-0.06	0.94	0
$\lvert\alpha_i - \alpha_j\rvert \times a$	0.51	1.66	0
$\lvert\gamma_i - \gamma_j\rvert$	-0.26	0.77	0

This table reports Cox proportional hazards results, augmented with measures of distance between skills.

have a higher match effect residual on average. There appears to be weak evidence of such a pattern.

I interpret the lack of evidence of sorting conditional on residual match effects to validate the exclusion restriction. If players anticipate match quality before partnering together, the set of extant matches is selected; then θ_i overestimates i's contribution to team production, and should instead be interpreted as a combination of team skill and the ability to search out good quality matches.

4.9 Survival analysis of teams

This subsection analyzes the determinants of team tenure. Table 10 reports parameter estimates from a Cox proportional hazards regression of skills on the team duration, the number of tournaments for which a given team ij plays together. Parameter estimates in levels and exponentials are given in the first two columns, followed by p-values. A unit increase in general skills input increases the likelihood of separation by 12%, while a unit

Table 12: Model fit verses omitted partner model

	Baseline		Partner omitted	
α	0.36***	0.34***	0.6***	0.57***
Ctrls.	-	x	-	x
Obs.	50035	50035	100070	100070
R^2	0.4	0.46	0.34	0.41
$\bar{R}^2$	0.38	0.43	0.33	0.39

This table compares the baseline model given by equation (6) to a model that omits the partner j fixed effect for each observation. To make complete use of the data I add a corresponding observation including the partner j fixed effect, but omitting that of the reference player i – for this reason the number of observations is doubled. For example, an observation of output y_{ijt} becomes a pair of observations $y_{it} = y_{jt}$. This simulates a dataset in which partner identity is unknown, but team outcomes are observed.

increase in team skills decreases the separation hazard by 17%. Teams relatively more intensive in team skill last longer, while general skill-intensive teams are more volatile. A team with a one unit higher residual match effect is 6% less likely to separate, indicating that partners remain together after observing high match quality.

Table 11 augments the previous regression with measures of difference between skills of teammates. The results of the initial regression remain approximately the same. A one unit greater difference between the general skill parameters makes a team 66% more likely to split, while a unit greater difference in team skill appears to have a negative effect on separation hazard rate in the order of 23%.

4.10 Comparison to omitted partner specification

The ATP dataset is atypical in that the exact identities of both team members are known. In contrast, typical labour force surveys do not contain information on coworkers.[27] This subsection compares the team production function given by equation (6) to a reduced form model in which the identity of the partner is unknown. Consider a team production specification omitting partner j's inputs.

$$y_{it} = \mu_2 + \tilde{a}(\alpha_i) + \tilde{\gamma}_i + X_{it}\tilde{\beta}_2 + \tilde{r}_{ijt} \tag{12}$$

[27] Deming (2017) notes that "wage returns [to interpersonal skills] cannot be identified without information about labor market frictions and about the skills of the other workers" (Deming, 2017, p.16).

This model compares a player i's solo performance to team performance, while failing to control to partner j attributes. Table 12 shows that this results in an overestimation to the transferability of general skills into team production; since general skills are correlated positively with partner's general skills, the reduced form model attributes greater magnitude to their input when partner's general skills are omitted. This model also underperforms compared to the baseline model by the metric of explanatory power.

4.11 Discussion

The above results show that there exists a fixed factor within individuals that explains performance in doubles tennis teams, above and beyond what can be predicted from the singles productivity of either partner. This is robust to relaxation of the log-additive functional form, and so is not an artifact of sub or superadditivity of general skill inputs. Such a skill factor may manifest in several ways. First, communication and interpersonal coordination come into play in doubles competition but are excluded from singles play, so the team skill factor absorbs deviations in these abilities from the mean.

Second, the team skill factor will pick up skills that are common to both singles and doubles play but employed in different intensities. For example, common understanding says that movement is less important in doubles play because there is less ground to cover per person, and that net play is more important because with greater court coverage, players are freer to approach the net. This implies that players with below average movement and above average net play will be found to have higher estimates of team skill. These players have a comparative advantage in team production compared to solo. However, making use of such skills requires coordination. For example, a player with poor lateral movement will not be able to exploit his comparative advantage if he does not coordinate with his partner; instead of dividing the court, he follows the ball into his partner's space and interferes with the return. Likewise, a player skilled in vollying will not feel safe approaching the net if he does not consider the response of his partner to provide cover to passing shots.

There is the question of idiosyncratic compatibility across teammates. This is picked up by the saturated match effect. Since match effects vary arbitrarily across matches even within a given player they are difficult to interpret in tangible terms. The Bryan brothers provide an illustrative example (although they are excluded from estimation for having not participated in singles tournaments during the sample years). Bob and Mike Bryan are a pair of American identical twins who have earned international fame as doubles champions. The twins have developed a system of code words to communicate strategies (Hansen 2013). Note that the baseline specification of the doubles model includes team tenure in number of tournaments. This nets out any average return to tenure that may accrue from teammates becoming accustomed to each other. However, teams that start with a higher degree of initial understanding for any reason will be identified as high match effect teams.

In addition to idiosyncratic properties of a match, the saturated match effects will absorb effects of specialization of roles within a team. For example, consider a player with powerful groundstrokes paired with a partner adept at net play. This team may be able to find success by gravitating the power player towards the back of the court and the finesse player near the net.[28] Now suppose the power player switches to a new partner who shares his general skill profile. There is now less scope for gains from specialization, and the new team may be less successful. In this case the former team will be identified as having a large match effect due to the compatibility of the partners; in this case the compatibility is not truly idiosyncratic, but systematically determined by the scope for specialization as determined by the production technology. This scope is probably small in tennis; players must alternate service games, and must alternate serve to either opponent. During a rally, each must be prepared to return. The set of tasks each must perform is largely similar.

The observed sorting patterns warrant some discussion. Early applications of AKM's two-way fixed effects decomposition on labour markets found insignificantly positive, or even negative assortative matching. This led to speculation that the wage determination function is misspecified, that limited mobility bias may be pushing down estimates of the correlation between worker and firm effects (Andrews et al. 2008), or that sorting truly is negatively assortative. Moreover, Eeckhout and Kircher (2011) show that sorting cannot be identified from wage data alone in frictional search environments. More recent applications have found evidence of positive assortative matching in the order of a 20% correlation between worker and firm effects (Card et al. 2013).

Although the current paper decomposes team output according to the same two-way fixed effects model, the setting is substantially different. First, prize money is awarded according to a fixed schedule increasing monotonically in performance, so Eeckhout and Kircher (2011)'s result does not apply.[29] More crucially, professional doubles tennis

[28] Anecdotally, this happens more frequently in women's doubles; in men's, partners tend to be equidistant from the net.

[29] The authors show that in an environment with wage bargaining, a given worker's wages do not in-

teams consist of two partners, in contrast to a job match from typical match employee-employer datasets where an observation consists of a worker and a firm. Rather than assuming there exists a log-additive fixed factor of wage determination attached to a firm, the current paper assumes only fixed factors of output determination attached to individual workers (and considers a relaxation of the log-additive structure). Since a team only consists of two players, no separability assumption between worker and firm/team is necessary – there are no other 'employees' attached to a match to consider. This reduces the space of interpretations as to what the fixed effect object captures compared to a firm fixed effect. The result is a stronger observed sorting pattern than seen in labour market applications. Moreover, partners sort positively assortatively according to general skill, an 'out of sample' measurement that comes from solo production rather than team play.

There is strong evidence of positive assortative matching within both skill classes. This is consistent with a setting in which utility is nontransferable between partners or there exist search frictions.[30] The former condition is decidedly the case in professional tennis: each partner receives half the prize money, and transfer payments are not made (see Appendix H for details). The issue of search frictions is less clear; information and coordination frictions are less likely to exist in higher leagues where events receive more publicity and players travel greater distances in order to compete. According to the current results, sorting as measured by the simple correlation tends to be stronger in the more elite leagues, where travel costs are lower relative to prize money, and information incompleteness is likely less intense. This suggests some role for search frictions.

Sorting across skill categories is not observed. This is not mechanical artifact of functional form specifications; players sort positively assortatively by composite skill θ, from which team skill γ is derived. The general skill α comes from a different sample, and so is not mechanically related to θ. The absence of cross-skill sorting results from partners sorting positively assortatively along own and team skill dimensions, combined with the orthogonality assumption between own and team skill. Note that the estimated correlations of cross-skill sorting vary across leagues (see figures 11, 12, and 13). A

crease monotonically with firm quality. To illustrate, suppose the production function is complementary in worker and firm inputs. Consider a worker employed at a firm to which she is perfectly matched, earning a wage w. If this worker were to move to a marginally better employer, she would be paid less, because the better firm would need to be compensated for forgoing the option value of employing a marginally better worker. In contrast, in the case of professional tennis, earnings are predetermined by a prize money schedule that is monotonically increasing in team performance.

[30] See Legros and Newman (2007), Bartolucci and Monzón (2014) respectively.

possible explanation is as follows. Teammates commit to playing together, which imposes certain restrictions on scheduling (the interviewees questioned in Appendix H claim that this is the case). Suppose that a player with a high value of general skill, but with a low value of team skill, is reluctant to commit because his priority is to play in the singles category. A player high in team skill may prefer to form a team with somebody who commits to making themselves available, which would reasonably be a partner with a comparative advantage in doubles play; that is, a partner with high team skills. Then players with high general skills end up partnered, as do players with high team skills. The results of the survival analysis regression corroborate this story, with high general skill players having higher hazard rates of team separation.

5 Conclusion

This paper has identified value-added of team skills net of general skills by comparing team and solitary productivity for a given worker. The identifying assumption is that general skills are required for both modes of production but team skills only for teamwork. Team skills explain 50% of across team output variation in men's professional tennis. This implies a comparative advantage in team production for some players. General skills account for 20% of across team output variation.

Returns to skills happen directly – a higher skilled tennis player wins more, and so earns more – as well as indirectly, through assignment, because a higher skilled player tends to match with a higher skilled partner. There is large variation in general skills; one standard deviation increase in general skill α yields 124% higher earnings per singles tournament. The direct return to general skills is smaller in doubles, providing around 42% higher earnings. However, this is augmented by indirect returns in the order of 21% due to being paired with a better partner. The corresponding direct and allocative returns to team skills in doubles are 66% and 34% respectively. A large portion of earnings for higher skilled players is due to being paired with more capable partners. Doubles partners sort positively assortatively within both skill dimensions, but exhibit no systematic sorting patterns across dimension.

These results are robust to relaxing the log-additivity assumption of the two-way fixed effects model of Abowd et al. (1999). A functional form test statistically fails to reject the log-additive specification in favour of a log-modular specification in general skill (see Appendix D). Moreover, the data fail to reject the two-way fixed effects model in favour of the more flexible match effects model. Analysis of the residual match effects indicates that players cannot anticipate the idiosyncratic component of team productivity before playing together as a team, implying that match quality is an experience good.

The results speak to optimal team size and the division of tasks within an organization. Some workers have a comparative advantage at team production, and some at solitary production. Firms allocating workers to tasks should plan accordingly. Benson et al. (2018) find that firms disproportionately emphasize general skills when making promotions to manager, resulting in costly misallocation of talents. The current paper provides a basis to compare solitary work to teamwork, so that workers can be allocated towards their comparative advantage.

Misattribution of contributions to team output may create an environment of uncertainty in which discrimination can persist. According to the current model a worker who performs poorly alone but systematically better in teams is a good team player; without this guidance such an individual may be misidentified as unproductive. Lack of acknowledgement for teamwork may contribute to labour market discrimination against women, who Kuhn and Villeval (2015) find to be more inclined to teamwork than men. Sarsons (2017) finds that women are not rewarded for their contribution to team production, while men are.

The identification contribution of this paper can be applied to other settings where team size varies. However, future applications will have to grapple with the problems of specialization of roles within a team and endogenous team size that are not of concern in the current setting.

References

Abowd, J. M., Creecy, R. H., and Kramarz, F. (2002). Computing person and firm effects using linked longitudinal employer-employee data. Longitudinal Employer-Household Dynamics Technical Papers 2002-06, Center for Economic Studies, U.S. Census Bureau.

Abowd, J. M., Kramarz, F., and Margolis, D. N. (1999). High wage workers and high wage firms. *Econometrica*, 67(2):251–333.

Andrews, M. J., Gill, L., Schank, T., and Upward, R. (2008). High wage workers and low wage firms: Negative assortative matching or limited mobility bias? *Journal of the Royal Statistical Society. Series A (Statistics in Society)*, 171(3):673–697.

Arcidiacono, P., Kinsler, J., and Price, J. (2017). Productivity spillovers in team production: Evidence from professional basketball. *Journal of Labor Economics*, 35(1):191–225.

Autor, D. H. and Dorn, D. (2013). The growth of low-skill service jobs and the polarization of the us labor market. *American Economic Review*, 103(5):1553–97.

Bartolucci, C. and Monzón, I. (2014). Frictions lead to sorting: a partnership model with on-the-match search.

Benson, A., Li, D., and Shue, K. (2018). Promotions and the peter principle. Technical report, National Bureau of Economic Research.

Brave, S. A., Butters, R. A., and Roberts, K. (2017). In search of the holy grail: Team chemistry and where to find it.

Buhai, I. S., Portela, M. A., Teulings, C. N., and van Vuuren, A. (2014). Returns to tenure or seniority? *Econometrica*, 82(2):705–730.

Card, D., Heining, J., and Kline, P. (2013). Workplace heterogeneity and the rise of west german wage inequality. *The Quarterly Journal of Economics*, 128(3):967–1015.

Chetty, R., Friedman, J. N., and Rockoff, J. E. (2014). Measuring the impacts of teachers I: Evaluating bias in teacher value-added estimates. *The American Economic Review*, 104(9):2593–2632.

Cortes, G. M., Jaimovich, N., and Siu, H. E. (2018). The "end of men" and rise of women in the high-skilled labor market. Technical report, National Bureau of Economic Research.

Deming, D. J. (2017). The growing importance of social skills in the labor market. *The Quarterly Journal of Economics*, 132(4):1593–1640.

Eeckhout, J. and Kircher, P. (2011). Identifying sorting—in theory. *The Review of Economic Studies*, 78(3):872–906.

Gibbons, S., Overman, H. G., and Pelkonen, P. (2014). Area disparities in Britain: Understanding the contribution of people vs. place through variance decompositions. *Oxford Bulletin of Economics and Statistics*, 76(5):745–763.

Gibson, A. D. (2010). The Association of Tennis Professionals: From player association to governing body. *The Journal of Applied Business and Economics*, 10(5):23.

Guryan, J., Kroft, K., and Notowidigdo, M. J. (2009). Peer effects in the workplace: Evidence from random groupings in professional golf tournaments. *American Economic Journal: Applied Economics*, 1(4):34–68.

Guvenen, F., Karahan, F., Ozkan, S., and Song, J. (2015). What do data on millions of us workers reveal about life-cycle earnings risk? Technical report, National Bureau of Economic Research.

Hansen, E. (2013). The secret language of tennis champions: How identical twins Mike and Bob Bryan serve science. *Nautilus Magazine*, 006.

Kambourov, G., Siow, A., Turner, L., et al. (2015). Relationship skills in the labor and marriage markets.

Kuhn, P. and Villeval, M. C. (2015). Are women more attracted to co-operation than men? *The Economic Journal*, 125(582):115–140.

Kuhn, P. and Weinberger, C. (2005). Leadership skills and wages. *Journal of Labor Economics*, 23(3):395–436.

Lazear, E. P. and Shaw, K. L. (2007). Personnel economics: The economist's view of human resources. *The Journal of Economic Perspectives*, 21(4):91–114.

Legros, P. and Newman, A. F. (2007). Beauty is a beast, frog is a prince: Assortative matching with nontransferabilities. *Econometrica*, 75(4):1073–1102.

Lise, J. and Postel-Vinay, F. (2016). Multidimensional skills, sorting, and human capital accumulation. *(Working Paper)*.

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

Ngai, L. R. and Petrongolo, B. (2017). Gender gaps and the rise of the service economy. *American Economic Journal: Macroeconomics*, 9(4):1–44.

Oettl, A. (2012). Reconceptualizing stars: Scientist helpfulness and peer performance. *Management Science*, 58(6):1122–1140.

Sarsons, H. (2017). Recognition for group work: Gender differences in academia. *American Economic Review*, 107(5):141–45.

Stevenson, B. (2010). Beyond the classroom: Using title ix to measure the return to high school sports. *The Review of Economics and Statistics*, 92(2):284–301.

A Identification

General skill, team skill, and idiosyncratic match effect indeces are defined up to a scale, so the goal is to identify α_i, γ_i, ψ_{ij}, for $i, j \in \mathcal{I}$, $i \neq j$, as well as the relative intensity of general skill in team production a.

A.1 Identification of the solo production model

Identification of general skill parameter vector α is trivial. Estimate equation (1) using the textbook within transformation. Impose a constraint $\sum_{i \in \mathcal{I}} \alpha_i = 0$ and recover the $I \times 1$ vector of general skills α.

A.2 Identification of the saturated match effects model

Because the general model of team production given by equation (2) places no restriction across matches that share a worker input i, such as ij and ik, observing a given worker i switch across multiple partners does nothing to aid identification. All that is required to estimate grand match effect Φ_{ij} is observation of ij's output. Impose a constraint $\sum_{ij \in \mathcal{M}} \Phi_{ij} = 0$ and recover the $M \times 1$ vector of general skills Φ.

A.3 Identification of the two-way fixed effects model

Abowd et al. (2002) show identification of the parameters of the two-way fixed effects model in a bipartite matching setting. This subsection will replicate their argument in the nonpartite setting and add the conditions needed to separately identify general and team skills, and will describe the circumstances under which they can be interpreted as measuring inputs into the team production function.

The two way fixed effects model given by equation (3) places restrictions on workers; namely, worker i's composite fixed effect θ_i remains fixed across teams. In order to identify the parameter vector θ the econometrician must observe worker mobility. To illustrate, abstract from time variables; then output of the match ij in time period t

is given by $y_{ijt} = \mu_2 + \theta_i + \theta_j + r_{ijt}$. Averaging across time yields $\bar{y}_{ij} = \mu_2 + \theta_i + \theta_j$, and differencing out $\bar{y}_{ik}$ yields $\theta_j - \theta_k$. Then $I - 1$ first differences are identified for a connected set of workers. Define a connected set as follows.

Definition 2 - *Connected set*

A set of workers $\mathcal{I}$ is a connected set *if for any workers i and k, $i, k \in \mathcal{I}$, $i \neq k$, there exists some sequence of matches $\{j, j+1\}_{j=i}^{k}$ such that $j, j+1 \in \mathcal{M}$.*

In other words, the set is connected if i has worked with someone, who has worked with someone – and so on – who has worked with k.

Assumption 3 - *Connectedness*

The set of workers $\mathcal{I}$ is connected.

For a set $\mathcal{I}$ containing I workers to be connected, the set of matches $\mathcal{M}$ must contain at least $I - 1$ elements. This implies $I - 1 \leq M \leq \frac{I(I-1)}{2}$. If the given data are not connected, typical practice is to consider only the largest connected set.

With $I - 1$ first differences of the form $\theta_j - \theta_k$ identified, one more degree of freedom is required to identify I worker fixed effects in levels. Most applications normalize $\theta_i = 0$ for some arbitrary θ_i; the current paper imposes the constraint $\sum_{i \in \mathcal{I}} \theta_i = 0$.

Unbiasedness of the estimate $\hat{\theta}$ requires the residual term $\tilde{r}$ from equation (3) be uncorrelated with observables. Unlike the saturated fixed effects model, the two-way fixed effects model relegates the idiosyncratic portion of the match effect (given by ψ) to the residual. This requires the following assumption.

Assumption 4 - *Independence of assignment from idiosyncratic match effects*

The matching function J can depend arbitrarily on all individual characteristics of i and j, but is independent of the idiosyncratic component of the match effect given by ψ_{ij}, and the error term ε_{ijt}.

This says that idiosyncratic match effects ψ_{ij} do not determine the assignment of matches. It is the identifying assumption used in the wage decomposition literature, and has been informally tested in that setting by Card et al. (2013). The current paper provides evidence consistent with assumption 4 in subsection 4.8.

Once θ and α are identified, their correlation yields the general skills transferability parameter a, and the difference between θ and $a\alpha$ gives the vector of team skills γ. Note that all elements of α and θ are deviations from the mean. Then when solving the system of I equations implied by equation (5), the elements of the residual vector γ are also in terms of mean deviations (a nonzero mean of θ or α would imply a nonzero mean of γ).

B Literature review

A large literature has investigated peer effects and spillovers. Typically this research implicitly assumes that it is the more productive workers who have greater capacity to improve their teammates' performance, skills being summarized by a single dimension. For example, Mas and Moretti (2009) find that grocery store clerks work faster when a productive coworker is in their line of sight; on the other hand, Guryan et al. (2009) find that the ability of a randomly paired partner does not have a large effect on golf performance. Teamwork ability overlaps conceptually with spillovers: a worker who cooperates successfully with partners may be said to have a large peer effect. The key distinction in the current paper is that teamwork skills vary independently from separable skills; the most productive solitary worker may not be the most effective teammate. An illustrative example is that of an introverted software engineer, who may be productive at coding, but is unable to coordinate with team members. An extreme case would be that of a star salesman who excels according to individual metrics by poaching commissions from colleagues. Managers concerned with team or firm level output would be unlikely to favour either such worker.

A small recent literature investigates teamwork ability as a distinct object from solitary ability. Arcidiacono et al. (2017) identify spillovers in NBA teams. Players whose presence during a given possession makes their teammates more likely to score are identified as having a large spillover effect, while players who themselves are more likely to score have a large own effect.[31] Observation of individual-level, within-game outcomes – namely the identity of the scorer – gives power to identify own and spillover effects separately. This requires general skills and spillover skills to be defined such that general skill contributes only to one's own scoring probability, and spillover skill contributes only to the scoring probability of another. The authors find evidence of two distinct factors. Even though it explains a substantial portion of variation in scoring outcomes, the spillover dimension is not rewarded in terms of salary, which in the NBA is determined by contracts based on a player's perceived value.

Oettl (2012) constructs a measure of helpfulness among academic immunology researchers, finding that coauthors of helpful scientists experience a decrease in research quality following the premature death of their peer. The author observes direct evidence of collaboration: acknowledgments by scientists of assistance from peers in research pub-

[31] The authors also estimate defensive ability from the probability that one's opponents do not score during a defensive possession.

lications. Scientists who receive many acknowledgments are labeled as helpful. Following the untimely death of a helpful scientist, coauthors experience a decrease in output quality, as measured by citations. There is no decrease in output quantity following such a death, nor is there any effect of a coauthor's death when that coauthor is not categorized as helpful.

The aforementioned papers speak to a division of tasks that may be endogenous to the team setting. For example, tenured professors may be free to focus on helping colleagues once their job security is no longer dependent on productivity attributed to their authorship. In this case there may be only one dimension of skill, but a division of tasks that produces multiple types of outcomes.

Brave et al. (2017) infer a measure of teamwork among Major League Baseball players whose teams win more frequently than the standard measure of wins above replacement (fWAR) would imply. The fWAR measure itself is a weighted average of observable metrics such as on-base average, each compared to the league average. Because most interactions in baseball are discrete, the primary channel for teamwork is through indirect spillovers – a player may perform differently when in the presence of certain teammates, although they are not physically interacting.

The current paper complements the above literature and contributes in the following ways. First, singles tennis performance provides a signal of player ability outside the context of any team. The individual productivity metrics from the papers mentioned above are based on individual-level outcomes measured within a team environment. To the extent that these outcomes are influenced by other team members they are endogenous to team spillovers and the team-level optimization problem. A productivity measurement coming from outside the team environment is uncontaminated by team influence. As in Oettl (2012) and Brave et al. (2017), outcomes measured at the team level ensure that individual qualities are in units of value-added; unlike these, the current paper measures value-added in 2015 US dollars. This is possible because of the prize schedule of professional tennis increases monotonically in team performance.

C Estimation

Several features of estimation require discussion. First, the large number of fixed effects makes matrix inversion computationally intensive. This is especially true for the saturated match effects model, which includes a fixed effect for each extant team. Therefore OLS solutions are calculated using the iterative conjugate gradient method (see Abowd et al. 2002). This returns the exact OLS solutions. Second, the different factor intensities in the team and solo production functions imply a nonlinear constraint between the equations that are to be estimated.[32] This requires either nonlinear estimation linear estimation in two stages.

This section outlines two equivalent two-stage linear estimators. The first is estimation of equations (1) and (6) by seemingly unrelated regression to yield the vectors of individual and compound skill α and θ, and once these are recovered, estimate the relative intensity of individual skill in team production a as well as the residual vector of team skill γ using auxiliary regression specification (5). Alternatively, estimate equation (1) by OLS to recover $\hat{\alpha}$, and include these estimates in a second stage regression estimation of equation (6). The latter method extends naturally to the nonlinear team production function given by equation (29), while the former does not.

C.1 Seemingly unrelated regression and auxiliary estimation of team skill

First collect the solo regression equations given by (1) in matrix form. Recall that there are I individuals observed over T time periods. A balanced panel has $I \times T$ solo production observations, and an unbalanced panel $N_1 < I \times T$. Letting D_1 be a $N_1 \times I$ design matrix of dummy variable indicators for player identity in solo production observations and X_1 a $N_1 \times k_1$ matrix of observables, the N_1 equations implied by equation (1) can be rewritten as

$$Y_1 = \mu_1 + D_1\alpha + X_1\beta_1 + \varepsilon_1 \tag{13}$$

where Y_1 is the $N_1 \times 1$ vector of observed output and α the $N_1 \times 1$ parameter vector of general skill fixed effects (a typical element being given by α_i).

Next, rewrite team production estimation equation (6) in matrix form. With a population of I workers there are $M^* \equiv \frac{I(I-1)}{2}$ possible teams to be formed. Since a given worker can only participate in one team at a time, a minimum of $T^* = I - 1$ time periods are required to observe every team and a balanced panel of $M^* \times T$ observations is impossible. Consider instead a sample of $N_2 \leq M^* \times T$ team production observations. The $N_2 \times 1$ output vector Y_2 can then be written as a function of the $N_2 \times I$ design matrix for the arbitrarily chosen first partner D_2 and that for the second partner F, as well as the $N_2 \times k_2$ matrix of observables X_2. This collects the N_2 equations implied by

[32] The relative intensity of general skill into team production is given by a; see solo production function (1) and team production (3).

(6) into the following system:

$$Y_2 = \mu_2 + DF_2\theta + X_2\beta_2 + \tilde{r}_2 \tag{14}$$

where $DF_2 \equiv D_2 + F$, imposing consistency.[33]

Now stack equations (13) and (14) into the following system of equations:

$$\begin{aligned}
\begin{bmatrix} Y_1 \\ Y_2 \end{bmatrix} &= \begin{bmatrix} \iota_{N_1} & 0 \\ 0 & \iota_{N_2} \end{bmatrix} \begin{bmatrix} \mu_1 \\ \mu_2 \end{bmatrix} + \begin{bmatrix} D_1 & 0 \\ 0 & DF_2 \end{bmatrix} \begin{bmatrix} \alpha \\ \theta \end{bmatrix} + \begin{bmatrix} X_1 & 0 \\ 0 & X_2 \end{bmatrix} \begin{bmatrix} \beta_1 \\ \beta_2 \end{bmatrix} + \begin{bmatrix} \varepsilon_1 \\ \tilde{r}_2 \end{bmatrix} \\
&= \begin{bmatrix} \iota_{N_1} & 0 & D_1 & 0 & X_1 & 0 \\ 0 & \iota_{N_2} & 0 & DF_2 & 0 & X_2 \end{bmatrix} \begin{bmatrix} \mu_1 \\ \mu_2 \\ \alpha \\ \theta \\ \beta_1 \\ \beta_2 \end{bmatrix} + \begin{bmatrix} \varepsilon_1 \\ \tilde{r}_2 \end{bmatrix} \qquad (15) \\
Y &= Z\xi + \varepsilon
\end{aligned}$$

where ι_n is a vector of ones of length I and the zeros are commutable zero matrixes. which can be estimated by including a set of linear constraints on the α and θ parameter vectors as follows.[34] Without loss of generality constrain the parameter vectors α and θ each to sum to zero. The system of constraints then appears as

$$C\xi = \begin{bmatrix} 0 & 0 & \iota'_{N_1} & 0 & 0 & 0 \\ 0 & 0 & 0 & \iota'_{N_2} & 0 & 0 \end{bmatrix} \xi = 0$$

which combined with (15) yields the constrained estimator

$$\begin{bmatrix} \hat{\xi} \\ \lambda \end{bmatrix} = \begin{bmatrix} 2Z'Z & C' \\ C & 0 \end{bmatrix}^{-1} \begin{bmatrix} Y \\ 0 \end{bmatrix} \tag{16}$$

[33] This constraint is necessary because no matter how workers are arranged into design matrixes D_2 and F – each with rows summing to one – some individual must necessarily fall into both matrixes unless the network structure is bipartite. To see this, consider that a non-bipartite network must contain an odd cycle. Without loss of generality number workers in this cycle by $\{1, ..., I\}$. Suppose that all odd numbered workers fall into design matrix D_2, and all even numbered workers into F. But by nature of being a cycle, workers 1 and I are observed working as a team. For this observation both workers, being odd numbered, should fall into D_2; but then the corresponding row of D_2 does not sum to one. This contradicts the supposition, completing the proof by contradiction.

[34] Alternatively drop one column each of the D_1 and DF_2 matrixes; the μ parameters will be estimated as the corresponding individuals' fixed effects, and the elements of α and θ will be relative to this baseline. Most of the wage determination literature uses this approach, but in the present case the levels of μ are of interest, since the correspond to the mean output values of solo and team production instances.

where λ is a 2×1 vector of Langranian multipliers. In the case that D_1 and DF_2 are too large to allow feasible inversion of the matrix $Z'Z$, obtain the OLS solutions by using the conjugate gradient method.

Now use the estimates $\hat{\alpha}$ and $\hat{\theta}$ contained in $\hat{\xi}$ to estimate the team skill parameter vector $\hat{\gamma}$ using the specification given by equation (5). Collecting the I equations in matrix form and allowing for measurement error yields

$$\begin{aligned} \hat{\theta} &= a\hat{\alpha} + \gamma \qquad (17) \\ \theta + u_\theta &= a(\alpha + u_\alpha) + \gamma \end{aligned}$$

where measurement error is given by u_v for $v = \alpha, \theta$. Recall that the parameter a is the relative intensity of general skill in team production. Team skill γ is given by the residual: it is the component of team value-added not explained by general skills. This implies it is orthogonal to general skill by construction.

There is measurement error in both the dependent and independent variables. Moreover, the error terms may be correlated. The estimator of relative general skill intensity is

$$\hat{a} = \frac{\mathrm{cov}(\hat{\alpha}, \hat{\theta})}{\mathrm{var}(\hat{\alpha})} = \frac{\mathrm{cov}(\alpha, \theta) + \mathrm{cov}(u_\alpha, u_\theta)}{\mathrm{var}(\alpha) + \mathrm{var}(u_\alpha)} \qquad (18)$$

where the classic errors-in-variables assumptions $\alpha \perp u_\alpha$, $\theta \perp u_\theta$ are augmented by a pair of cross-orthogonality assumptions $\alpha \perp u_\theta$ and $\theta \perp u_\alpha$.

The primary concern is bias due to $\mathrm{cov}(u_\alpha, u_\theta) \neq 0$. For example, when solo and team production are both observed in a given time period for a given individual, a common shock may apply to both observations, biasing the correlation upwards.[35] Conversely, one mode of production may take precedence over the other in such a case, biasing the correlation downwards. As a robustness check, omit such observations. Subsection 4.5 indicates that such bias is small. Second, measurement error in $\hat{\alpha}$ can bias the estimate of a downwards. To occlude both concerns, limit the sample to individuals observed for at least $\bar{T}$ periods in either production category. The baseline estimates are for $\bar{T} = 20$.

C.2 Two stage estimation

This subsection details the two-step estimation procedure. First recover general skill parameters $\hat{\alpha}$ from an estimation of equation (13). Then include these estimated param-

[35] The scope for such shocks in the application to professional tennis is limited. For example, an injury may cause a player to withdraw from both categories and accept a smaller prize; but this is not measurement error, since by virtue of failing to complete the tournament, the player is identified as having a lower skill index.

eters as regressors in a team output regression. The vector of estimated parameters now treated as data is given by the $N_2 \times 1$ vector $\hat{A}_2 \equiv DF_2\hat{\alpha}$. Rewrite team specification (6) accounting for measurement error as

$$y_{ijt} = \mu_2 + a(\hat{\alpha}_i + \hat{\alpha}_j) + \gamma_i + \gamma_j + X_{ijt}\beta_2 + \hat{r}_{ijt} \tag{19}$$

where $\hat{r}_{ijt} = r_{ijt} - a(u_i + u_j)$. To the extent that the first stage estimates $\hat{\alpha}$ differ from the true parameters, the measurement error shows up in the error term. In matrix form:

$$\begin{aligned} Y_2 &= \mu_2 + DF_2(a\hat{\alpha} + \gamma) + X_2\beta_2 + \hat{r}_2 \\ &= \mu_2 + \hat{A}_2 a + DF_2\gamma + X_2\beta_2 + \hat{r}_2 \end{aligned} \tag{20}$$

where a typical element of the the vector $\hat{A}_2 \equiv DF_2\hat{\alpha} = DF_2(\alpha + u_\alpha)$ is the sum of team ij's general skills $\hat{\alpha}_i + \hat{\alpha}_j$. Note that by this definition $P_{DF_2}\hat{A}_2 = \hat{A}_2$ (where $P_X \equiv X(X'X)^{-1}X'$ is the projection matrix); in other words, team-summed general skills vector $\hat{A}_2$ exists in the 'team space' spanned by DF_2.

A typical element of the residual $\hat{r}_2 \equiv r_2 - aDF_2u_\alpha$ is given by $\hat{r}_{ijt}$. With measurement error from the first stage in the residual, there is concern of bias in estimates of a and γ; β_2 is estimated without bias according to the previous assumptions. Then the estimator of relative general skill intensity is

$$\hat{\hat{a}} = \frac{\text{cov}(M_\iota M_{X_2}\hat{A}_2, M_\iota M_{X_2}Y_2)}{\text{var}(M_\iota M_{X_2}\hat{A}_2)} \tag{21}$$

$$= \frac{\text{cov}(M_\iota M_{X_2}DF_2\alpha, M_\iota M_{X_2}Y_2) + \text{cov}(M_\iota M_{X_2}DF_2u_\alpha, M_\iota M_{X_2}Y_2)}{\text{var}(M_\iota M_{X_2}DF_2\alpha) + \text{var}(M_\iota M_{X_2}DF_2u_\alpha)} \tag{22}$$

where $M_X \equiv 1 - P_X$ is the annihilation matrix for a given regressor X. In the latter expression, the second term in the numerator is equivalent to the $\text{cov}(u_\alpha, u_\theta)$ term in equation (18): the covariance between measurement error in own- and team skill. As stated above, omitting joint-category tournaments as a robustness check addresses concerns over this source of bias. The second term in the denominator approaches zero as the panel size within an individual grows large; for this reason the baseline specification retains individuals observed for a minimum of $\bar{T} = 20$.

C.3 Estimation of saturated match effects model

The saturated match effects model captures all time-invariant across-team variation. Consider a design matrix G that includes a dummy variable corresponding to each of the M extant matches in set $\mathcal{M}$. The dimensions of G are $N_2 \times M$. Then the N_2

equations implied by the saturated match effects production function (given by equation 2) are written in matrix form as

$$Y_2 = \mu_2 + G\Phi + X_2\beta_2 + \varepsilon \tag{23}$$

which requires that constraining the match effects to sum to zero; that is, $\sum_{ij\in\mathcal{M}} \Phi_{ij} = 0$ (alternatively drop one of the columns of G and use the corresponding team as the base category). Denote this constraint in matrix form as

$$C_G\xi_G = \begin{bmatrix} 0 & \iota_M & 0 \end{bmatrix} \begin{bmatrix} \mu_2 \\ \Phi \\ \beta_2 \end{bmatrix} = 0 \tag{24}$$

which implies that the constrained system of equations is given by

$$\begin{bmatrix} \hat{\xi}_G \\ \lambda_G \end{bmatrix} = \begin{bmatrix} 2\begin{bmatrix} \iota_{N_2} & G & X_2 \end{bmatrix} \begin{bmatrix} \iota_{N_2} \\ G \\ X_2 \end{bmatrix} & C'_G \\ C_G & 0 \end{bmatrix}^{-1} \begin{bmatrix} Y_2 \\ 0 \end{bmatrix} \tag{25}$$

where λ_G is the Lagrangian multiplier and the 0 terms are commutable zero matrixes.

D Nonlinear model

D.1 Nonlinear team production

The model function characterized by equations (1), (3), and (5) imply that team production is log-linear in inputs. Consider an alternative team production function

$$Y_{ijt} = A_2 F(H_i, H_j)(W_i W_j)^{\tilde{\gamma}} R_{ijt} \tag{26}$$

that allows for more flexibility in how individual skills interact. For example, defining

$$F = (H_i H_j)^{\tilde{a}\bar{\alpha}} \exp(b[(\ln H_i)^2 + (\ln H_j)^2]) \exp(c \ln H_i \ln H_j) \tag{27}$$

where $\tilde{b} = \frac{b}{\bar{\alpha}^2}$, $\tilde{c} = \frac{c}{\bar{\alpha}^2}$ yields log output function

$$y_{ijt} = \mu_2 + \tilde{a}(\alpha_i + \alpha_j) + \tilde{b}(\alpha_i^2 + \alpha_j^2) + \tilde{c}\alpha_i\alpha_j + \tilde{\gamma}_i + \tilde{\gamma}_j + X_{ijt}\beta_2 + \tilde{r}_{ijt} \tag{28}$$

which can be estimated by OLS once the $I \times 1$ individual skill parameter vector α is recovered from estimation of (1). A variety of specifications of $F(.,.)$ can be tested against the linear model (6). Functional form (27) provides a flexible framework to capture nonlinearities in general skill inputs (via the squared terms), as well as interactivity. Interactivity is particularly interesting because it violates the log-additivity assumption of the two-way fixed effects model: a significant positive (negative) estimate of $\tilde{c}$ indicates that there is log-supermodularity (submodularity) in inputs. By capturing nonlinear effects of scaling, inclusion of the squared terms occludes concern over spurious identification of complementarity or substitutability due to misspecification of the linear inputs.[36]

D.2 Identification of nonlinear model

Once the general skill parameter vector α is recovered from one-way fixed effects estimation of equation (1), plug α_i, α_j into the system of equations given by (29). Any combination of α_i, α_j can be treated as data with the caveat that they are measured with error.[37] Taking the log of equation (26) yields the general formulation

$$y_{ijt} = \mu_2 + f(H_i, H_j) + \tilde{\gamma}_i + \tilde{\gamma}_j + X_{ijt}\beta_2 + \tilde{r}_{ijt} \tag{29}$$

where $\ln F(.,.) \equiv f(.,.)$, from which team skill parameter vector $\tilde{\gamma}$ can be recovered from a two-way fixed effects decomposition so long as $f(.,.)$ is estimable. Allowing a flexible functional form in general skill inputs H reduces concern that the individual fixed factors in team skill vector $\tilde{\gamma}$ are artifacts misspecifying $f(.,.)$ in the additive case.

Taking the log of functional form (27) yields

$$f(H_i, H_j) = \tilde{a}(\alpha_i + \alpha_j) + \tilde{b}(\alpha_i^2 + \alpha_j^2) + \tilde{c}\alpha_i\alpha_j \tag{30}$$

where $\alpha_i \equiv \alpha \ln H_i$. Since the vector of general skill parameters α is estimated from the first stage, $\tilde{a}$, $\tilde{b}$, and $\tilde{c}$ are all identified, as are the remaining parameters of (29).

Note that in the nonlinear case – that is, when $f(x, y) \neq c(x + y)$, where c is some constant – team skill parameter vector $\tilde{\gamma}$ cannot be recovered from an auxiliary regression after general skills α and composite skills θ are recovered from simultaneous or separate regression. This is because in the nonlinear case the separability assumption of two-way fixed effects decomposition is misspecified, so the estimates of θ would be invalid under this assumption.

[36] For example, if the true model parameters are $\tilde{b} = \bar{b}$, $\tilde{c} = 0$, and workers sort perfectly positively-assortatively, omission of the squared terms will produce estimate $\hat{\tilde{c}} = 2\bar{b}$. Similarly, weaker forms of positive-assortative matching would produce spurious significant estimates of $\tilde{c}$.

[37] For a discussion of measurement error see the Estimation section.

Figure 17: Joint distribution of γ, all tournaments – nonlinear

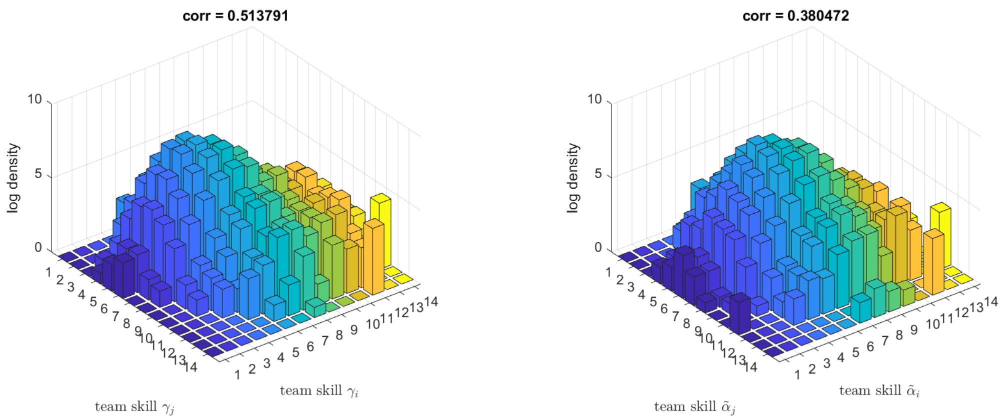

D.3 Two stage estimation of nonlinear model

Estimation of the nonlinear model in two steps extends naturally from the two-step procedure. First recover $\hat{\alpha}$ by estimating (13) by OLS. Then, accounting for measurement error, rewrite nonlinear output function (29) as

$$Y_2 = \mu_2 + DF_2(\tilde{a}\hat{\alpha} + \tilde{b}\hat{\alpha}^2) + \tilde{c}(D_2\hat{\alpha} \circ F_2\hat{\alpha}) + DF_2\tilde{\gamma} + X_{ijt}\beta_2 + \tilde{r} \tag{31}$$

where $\hat{\alpha}_i = \alpha_i + u_i$ and $\circ$ denotes the Hadamard product. The squared and interactive terms may introduce multiplicative measurement error. However, the estimates of team skill allowing for nonlinearities, given by $\tilde{\gamma}$, closely resembles that of the baseline team skill parameter vector γ. This suggests the multiplicative measurement error does not cause substantial bias.

D.4 Results of nonlinear estimation

Figure presents scatter plots of team skill as estimated by the linear and nonlinear models respectively. The sorting pattern is largely similar, although the correlation between skills of partners is smaller in the nonlinear model.

Table 2 shows that the nonlinear extension adds little explanatory power compared to the baseline linear model.

E Alternative variance decomposition measurements

This section details the measures of variance decomposition used in the results section and introduces several alternatives. These are the raw variance share (RVS), the correlated variance share (CVS), and the uncorrelated variance share (UVS). Compare these with the balanced variance share (BVS), detailed in section 4.

The first four methods are as follows. The raw variance share (RVS) for a given group of regressors is the of R^2 a regression containing only that group of regressors.[38] The correlated variance share (CVS) is the variance of one group of predicted coefficients times regressors $X_k\hat{\beta}_k$ relative to the variance of the dependent variable. The uncorrelated variance share (UVS) is the R^2 of the full regression specification less the R^2 of a regression that omits the group of regressors in question; one can think of it as the marginal contribution to explanatory power of that regressor group. The balanced variance share (BVS) is the variance of one group of regressors, summed with the covariances between that group and the other regressors, all relative to the variance of the dependent variable. The BVS sum to the R^2 of the full regression model. See Gibbons et al. (2014) for a full discussion.

The final two columns of Table (6) give the simple correlation between a predicted regressor group and the dependent variable, and the standard deviation of that predicted group, respectively. These are the measurements of variance preferred by AKM.

F Professional tennis circuits

Section 3 presents summary statistics and a broad description of the data. This section presents some additional information on how the different circuits mentioned in section 3 are related. Note that the circuits share a common ranking system; ranking points won in any circuit determine eligibility in any other.

Table 18 shows summary statistics by circuit. Top prizes vary substantially, and higher level leagues have greater variation in prize money outcomes. Players in elite leagues tend to be older; players who start in the lower leagues work their way up. Higher levels of experience in the more elite leagues attests to this.

Table 19 describes the different levels of tournament. There are many more tournaments per year in lower circuits, and tournaments are smaller.

[38] For example, the R^2 of the regression $Y^I = D\alpha + \varepsilon$ is 0.42, as seen in the first row and column of table (6).

Table 13: Variance decomposition for singles regression

	RVS	CVS	UVS	BVS	$\rho(.,Y)$	Std(.)	F	p
α	0	0	0	0	0	0	0	1.0
experience	0.38	0.32	0.3	0.34	0.6	1.4	41.91	0
age	0.09	0.01	0	0.02	0.21	0.25	0	1.0
year	0.09	0.01	0	0.02	0.21	0.25	0	1.0
month	0.57	0.02	0	0	0.02	0.36	50.05	0
tourney	0.3	0.01	0.01	0	0.06	0.18	131.13	0

Table 14: Variance decomposition for doubles regression – linear

	RVS	CVS	UVS	BVS	$\rho(.,Y)$	Std(.)	F	p
α	0	0	0	0	0	0	0	1.0
γ	0.03	0.13	0	0.1	0.29	0.81	0	1.0
experience	0.4	0.31	0.17	0.28	0.51	1.26	7.47	0
tenure	0.15	0.02	0	0.03	0.21	0.32	1.36	0
age	0.03	0	0	0.01	0.13	0.11	0.81	0.97
year	0.18	0.01	0	0	0.06	0.19	1.14	0.07
month	0.42	0.04	0.01	0.03	0.16	0.45	190.22	0
tourney	0.22	0	0	0	0.03	0.15	24.76	0

Table 15: Variance decomposition for doubles regression – nonlinear

	RVS	CVS	UVS	BVS	$\rho(.,Y)$	Std(.)	F	p
α	0	0	0	0	0	0	0	1.0
α^2	0.03	0.06	0	0.07	0.29	0.56	0	1.0
$\alpha_i \times \alpha_j$	1.24	0.07	0	0.1	0.39	0.6	0	1.0
γ	0.48	0	0	0	-0.24	0.02	1.26	0.26
experience	0.4	0.19	0.12	0.21	0.49	0.99	5.39	0
tenure	0.15	0.02	0	0.03	0.21	0.32	1.36	0
age	0.03	0	0	0.01	0.13	0.11	0.81	0.97
year	0.18	0.01	0	0	0.06	0.19	1.14	0.07
month	0.42	0.04	0.01	0.03	0.16	0.45	189.89	0
cons.	0.22	0	0	0	0.03	0.15	24.7	0

Table 16: Variance decomposition for singles regression – including tournament controls

	RVS	CVS	UVS	BVS	$\rho(., Y)$	Std(.)	F	p
α	0	0	0	0	0	0	0	1.0
experience	0.38	0.44	0.22	0.39	0.59	1.63	40.48	0
age	0.09	0.02	0	0.03	0.21	0.39	0	1.0
year	0.09	0.02	0	0.03	0.21	0.39	0	1.0
month	0.57	0.06	0	0	0.02	0.58	147.09	0
tourney	0.3	0	0	0	0.03	0.14	22.13	0
cons.	0.28	0.21	0.15	0.08	0.19	1.12	25.12	0

Table 17: Variance decomposition for doubles regression – including tournament controls

	RVS	CVS	UVS	BVS	$\rho(., Y)$	Std(.)	F	p
α	0	0	0	0	0	0	0	1.0
γ	0.03	0.02	0	0.04	0.29	0.35	0	1.0
experience	0.4	0.15	0.07	0.18	0.46	0.88	3.15	0
tenure	0.15	0.02	0	-0.02	-0.12	0.29	1.37	0
age	0.03	0	0	0	0.08	0.09	0.68	1.0
year	0.18	0.01	0	0	-0.02	0.18	1.08	0.18
month	0.42	0.08	0.01	0.05	0.16	0.66	212.38	0
tourney	0.22	0	0	0	-0.01	0.14	4.57	0
cons.	0.4	0.19	0.07	0.27	0.61	1.0	2.85	0

Table 20 shows the overlap in player participation across tournament type. Most of the players who compete in the ATP500 level tournaments participate also in higher level tournaments; the fraction is substantially small for ATP250 level tournaments, and smaller still for Challenger and Futures tournaments. However the converse is not true: most players who participate in higher level tournaments have played in the lower level circuits during the sample period. This suggests and 'up or out' practice: a player either works his way up to the elite leagues, or leaves the sport.

G Rules of tennis

A tennis court is 24 meters long (78 feet) with a net separating play areas in the middle. Each player or team remains on one side of the net until a switchover occurs. All standard tennis courts are marked with doubles lanes to their right and left extremities, which are inactive during singles play. These widen the play area from 8 to 11 meters (27 to 36 feet). Players or teams strike the tennis ball with their rackets to propel it to the other side of the net, where it must land within the play area of the court in order to be considered; landing outside of the court (or falling on one's own side of the court or into the net) counts as an error, and the striker loses a point. Landing inside the play area on the other side of the net keeps the ball in play and prompts the opponent to strike it, facing the same parameters as the original striker. The opponent must return the ball across the net before it has bounced twice; the opponent can also return the ball before it has touched the ground once (this is called a volley). If the opponent fails to strike the ball or commits an error, a point is awarded to the original striker. Points are counted at discrete intervals in the space $\{0, 15, 30, 40\}$, and exceeding this space while remaining at least two intervals ahead of the opponent concludes a game. Once the server has won or lost the game, it is the opponent's turn to serve. If the preceding game is odd-numbered within the set, the players change ends; if it is even they remain on their current end. Games add up into sets, and sets into matches.

A tennis match is won by winning two of three sets (or three of five in men's grand slam singles). A set is won by winning at least six games while remaining at least two games ahead of one's opponent.[39] Each game is associated with service by one player, whether playing alone or as part of a doubles team. The server alternates sides every time a point is scored, serving from the ad or deuce court, which are to the server's left

[39] In some tournaments a tiebreak is played instead of sustained play until one player leads by two sets.

Table 18: Summary Statistics by Circuit

All tournaments

	Singles prize	Doubles prize	Age	Experience	Year	Month
mean	4648	780	24	65	2012	6
p50	292	141	23	54	2012	6
sd	42838	4238	4	48	2	3
min	0	0	14	1	2009	1
max	3800000	285000	42	249	2015	12
count	172779	50603	273985	273985	223382	223382

ATP World Tour + ITF Grand Slams

	Singles prize	Doubles prize	Age	Experience	Year	Month
mean	24620	5703	27	90	2012	6
p50	7309	2902	27	86	2012	6
sd	101949	12020	4	54	2	3
min	0	0	14	1	2009	1
max	3800000	285000	42	239	2015	12
count	29117	5257	34374	34374	34374	34374

ATP Challenger Series

	Singles prize	Doubles prize	Age	Experience	Year	Month
mean	1125	451	25	79	2012	7
p50	584	285	25	71	2012	7
sd	1981	507	4	52	2	3
min	0	0	14	1	2009	1
max	28105	4284	42	249	2015	12
count	43147	12121	55268	55268	55268	55268

ITF Futures Events

	Singles prize	Doubles prize	Age	Experience	Year	Month
mean	315	112	23	52	2012	7
p50	191	93	22	42	2012	7
sd	345	113	3	40	2	3
min	0	0	14	1	2009	1
max	2198	522	42	241	2015	12
count	94819	31781	126600	126600	126600	126600

Table 19: Pro tennis tournament classes

Tournament class	Events / year	Org.	Main draw size		Median top prize (2015 USD)		Unit
			Singles	Doubles	Singles	Doubles	
Olympic Games	$\frac{1}{4}$	IOC	64	32	-	-	Nation
Davis Cup	1	ITF	2	2	-	-	Nation
Grand Slam	4	ITF	128	64	3,093,000	225,000	Player
Masters	9	ATP	128	64	867,000	113,000	Player
ATP500	13	ATP	32-48	16-24	343,000	51,000	Player
ATP250	40	ATP	28-48	16	106,000	13,000	Player
Challenger	∼150	ATP	32	16	8,200	1,500	Player
Futures	∼600	ITF	32-48	16-24	1,440	310	Player

Prize money is in 2015 USD, per player (each member of a doubles team receives the amount listed). Since prize money varies across tournaments within a tournament category, this table reports the median.

Table 20: Player overlap across tournament classes

slam	atp1000	atp500	atp250	challenger	future
1	.78	.72	.35	.13	.06
.65	1	.61	.27	.1	.05
.71	.73	1	.32	.12	.06
.96	.9	.88	1	.29	.18
.97	.96	.92	.81	1	.42
.78	.76	.79	.86	.73	1

Each cell shows the fraction of players from the column league who play in the row league. Diagonal elements are necessarily equal to one, but the matrix is not symmetric. For example, only 6% of futures players are observed to play also in a slam (top right), while 78% of slam players are observed also in futures (bottom left).

and right hand side respectively. In singles, the returner can stand anywhere on the court but naturally positions oneself diagonally from the server in order to return. In doubles, returner choose a side to return from and cannot switch until the conclusion of a set, guaranteeing that all combinations of server to returner happen with regular frequency. The server stands on the baseline to serve, anywhere from the center of the court to the left or right end, and must serve by hitting the ball diagonally across the court into the opponent's service box. The server should not step across the baseline into the court as part of the service motion, but this rule is unevenly enforced. If the net stops the ball or the ball falls outside the service box, it is called a fault and the server has one more chance to serve – called a second serve – before a double fault occurs, which awards a point to the opponent. If the serve hits the net but continues into the service box, it is called a let and the server resumes the first serve.

To enforce the rules an umpire sits at an elevated position to the side of the court, overlooking play. Volunteer linesmen spot whether a ball has fallen into or outside the court, audibly calling outs. If a player disagrees with the call he can challenge it. This prompts review by a system of coordinated cameras that record a ball's position with a great deal of accuracy. A player gets a total of three unsuccessful challenges before being cut off, or an infinite number of successful challenges. There exists also a referee in addition to the umpire who may enter the court in order to enforce rules in exceptional cases.

In professional men's competition, coaches are disallowed from directly engaging players during play, and can usually be found in the stands. Players may communicate with coaches via eye contact or shouting. Coaches may debrief their player following a match. Coaches are not mandatory and some players play without one.

H Player and coach interviews

To supplement the data and provide context, myself and a colleague attended the 2015 Canadian Open to interview active tennis professionals willing to share their perspectives and experiences. In 2018 I returned to the tournament.[40] Together we interviewed

[40]The Canadian Open is a joint tournament between the ATP and the Women's Tennis Association (WTA), with men's events being held in Toronto in even-numbered years and Montreal in odd-numbered years, and women's events being held conversely. In the ATP tour the Canadian Open is a 1000-level tournament (one tier below a slam). We attended the Toronto events both times, and so got exposure to professionals from both tours.

several top doubles players, as well as the coaches of top singles and doubles players. Yaroslava Shvedova has 13 career doubles titles and a peak doubles ranking of number 3 in the world, as well as a singles title and a peak ranking of 25. Feliciano López has six career singles titles and four in doubles, with respective peak rankings of 9 and 12. Jiri Fencl coaches Lucie Hradecka, who was at the time partnered with fellow Czech Andrea Hlavackova. They have 22 and 25 doubles titles respectively, winning many of those together, and as a pair won the silver medal in the 2012 Olympic Games. Finally, Emmanuel Planque coaches top singles player Lucas Pouille, who has five singles titles and a career high singles ranking of number 10. All players have been active in both singles and doubles categories over the course of their careers. This appendix summarizes the pertinent points from our conversations with Ms Shvedova, Mr López, Mr Fencl, and Mr Planque. A great thanks to all for sharing their knowledge.

Singles versus doubles

Singles and doubles tennis are the same sport, emphasizes Mr López: they require the same set of skills. Ms Shvedova agrees, and adds that there is more strategy in doubles play compared to singles. Mr Planque points out that for top singles players, doubles competition is not a priority, but is a good opportunity to practice and stay active.

Doubles partners share prize money evenly

Although tournament organizers distribute prize money to winners in the form of personalized cheques, the economist's mind wanders to more complex contracts between partners. Both sources agree emphatically that players do not make side payments: prize money is shared between partners exactly how it is doled out by the tournaments, a 50-50 split. Any other arrangement would be difficult to contract on, and would likely be regarded as unsportsmanlike. The only caveat to this arrangement is that when a team has reached such a stage of a tournament as to qualify for prize money, and one partner withdraws due to injury or another concern, that partner is penalized and receives an attenuated payment, while the uninjured teammate receives full payment corresponding to the current round attained. Note that the penalty is imposed by tournament organizers; the withdrawing player receives a cheque with the penalty deducted, and no side payments are made.

Not just two players – a team

Players commit to the concept of a team as a unit. This necessarily indicates feelings of camaraderie and partnership, as well as mutual responsibility: to make oneself available for tournaments,[41] to approach team play competitively, to maintain a professional relationship with one's partner. Some teams last for many years, and a robust relationship can develop over time as teammates become familiar with each other's personalities and play styles. Following a tough loss, one player may not communicate with the partner for several days. Players who know each other well will understand that their relationship has not been strained, that this is a natural response to a loss; those less familiar may feel uncertain about their partner's state of mind.

Communication and mutual understanding are key. Viewers will note that in between every point – win or lose – doubles players acknowledge each other by giving five. This establishes a constant line of nonverbal communication in addition to explicit strategizing between points. Players are often seen covering their mouths while speaking before service so as to withold plans from the opponents. An alternative may be communicating in a language unfamiliar to the opponents, or in the exceptional case of the Bryan brothers, speaking in code words – according to Nautilus Magazine, the identical twins who have played hundreds of tournaments together have along the way developed a series of convenient shorthands for common strategies.

Team formation

Teams form based on inherent compatibility. Certain play styles may complement each other to varying extents. A left-handed partner – statistically a rarity – is seen as a competitive advantage, since opponents must face service from alternating orientations.

An unusual form of complementarity – if not in production than in deciding tournament eligibility – is national consanguinity. Certain tournaments, such as the Davis Cup and the Olympic Games, are centered around international rather than interpersonal competition, and so require a shared national origin. Mr Fencl notes that leading up to the summer Olympics, players have a particular preference for partnering with a countryman – naturally with a view to represent their country together at the games.

[41] During the off-season, teammates plan a rough schedule of which events to play throughout the course of the coming season.

Teammates train separately

Typically partners do not train together outside of competition.[42] Each player practices with their own entourage, usually including a coach, physical therapist, and hitting partners (although some players do not employ coaches). For the most part, teammates' shared play experience consists of their observable professional record.

Team dissolution

Players do not care per se if their individual talent levels diverge over the course of their partnership – they split only if they stop performing as a team. Divergent career priorities may also lead to a split: one player decides to slow down and play fewer doubles tournaments per year – a light schedule of around ten – while the partner prefers to maintain a full schedule of around twenty. Age difference is not a concern in itself, except to the extent that it may affect the previous point; but as shown in the Section 3, doubles careers are more conducive to age. Current success is a better predictor of expected team duration than age. Mr Fencl reports that team dissolution is usually a mutual agreement, but of course can be unilateral in certain cases.

Team reformation

Practical experience as a team is not limited to continuous partnerships; former partners who resume playing together pick up where they left off.

Prioritizing tournaments

Winning an Olympic medal changes your life. Winning a major is the next best thing, yielding a great deal of prize money and ranking points – which are used to gain admittance to tournaments and more prize money. The ATP World Tour Finals doubles category is regarded with equal esteem. Lower tier tournaments are frequent enough to be considered routine for the top players.

[42] A high profile exception is the team comprised of American identical twins Bob and Mike Bryan, former doubles world number ones and Olympic gold medal holders. The duo has released extensive footage of team training drills.

I Players and coaches

Players typically report one coach, but sometimes two. In these cases there is no indication that one coach specializes in singles or doubles. Anecdotally, singles specialists (who do not play enough doubles tournaments to be included in the estimation sample) may employ two coaches concurrently (see the example of Novak Djokovic below) and players active in both singles and doubles categories may employ a single coach (see the example of Lucie Hradecka in the Interviews Appendix in the main document). For the most part multiple players do not share a single concurrently, but in a small number of cases they do (more on this below).

To the extent that a coach improves a player's performance across both play categories, the coach's ability is captured in the general skill; if one or more coaches improve a player's performance differentially across categories, this is captured in the team skill (increasing the estimate if the coach improves performance in doubles, decreasing it if the coach improves performance more in singles). Furthermore, if the coach systematically helps the player to find partners well-suited to the player's skillset, this is captured in the team skill estimate.

I do not control for coach's ability, so all player skill estimates pick up also the abilities of the coaches. With panel data on coach identity a difference-in-difference approach could identify coach value-added; however, the existiting datset does not contain this information (see below). So a player's skill estimates should be interpreted as including their ability to find helpful coaches over the course of their career.

I.1 Coach data

Unlike the players, for whom I observe for each player a time series spanning the entire professional career, I have only cross-sectional data linking coaches to players.

Two variables that can vary but are reported as static are place of residence and coach. The latter is updated periodically; for example, as of 2016 when the data were scraped the coach for Novak Djokovic was reported as "Boris Becker and Marian Vajda". Djokovic has undergone a number of coaching changes since and his coach is currently reported to be "Marian Vajda". Since there is no retrospective data on coaching, the only information contained in my dataset is the identity of the coach as of the date of scraping.

For the 2016 cross-section I observe 594 unique coaches among the estimation sample of 1993 players. For the remaining players the data are either missing, or may indicate

Table 21: Players per coach for 2016

Players per coach	Number of coaches
1	530
2	49
3	12
4	1
5	1
6	1
Total	594

that the players do not employ full-time coaches (as Nick Kyrgios famously chooses not to). Higher level players – those who have reached more elite levels of tournament – are more likely to report a coach.

Table 21 reports that the number of players coached concurrently ranges from one to six, and reports the number of coaches who coach each number of players. The vast majority coach a single player, a small number two players, and an insignificant number more than this. This suggests a minor role for coaches managing many players concurrently to affect estimates, but since over 10% of players share a coach with another player, further investigation is warranted.

I.2 Are players with the same coaches in singles more likely to be matched together?

With the above data limitations in mind I attempt to answer this question directly. For 367 doubles teams in the estimation sample I observe the coach of both players, with 328 unique coaches. The two team mates share the same coach in 11 of the 367 teams.

Consider a set of players $\mathcal{P}$ containing P unique individuals. At a given point in time these players employ a set of coaches $\mathcal{C}$ containing $C \leq P$ unique coaches. Assume that each player employs one coach, but a coach can work for multiple players.

For each coach $c \in \mathcal{C}$, let the number of players coached be given by $p_c \geq 1$. Then for each coach c, the number of potential teams of two players who are both coached by coach c is given by the following equation:

$$M_c^* \equiv \frac{p_c \times (p_c - 1)}{2}$$

where the asterisk denotes latency; let the number of extant matches between players of coach c reported in the data be given by $M_c \leq M_c^*$.

Likewise, the total number of potential teams among all players $p \in \mathcal{P}$ is given by $M^* \equiv \frac{P\times(P-1)}{2}$. Therefore the fraction of players who share a coach in a fully connected set (in which every player is matched to every potential partner) is given by $\frac{\sum_{c\in\mathcal{C}} M_c^*}{M^*}$.

The information in Table 21 can be used to calculate this fraction for the sample of players in the data for whom the coach is observed (for the remaining players I am unable to distinguish between nonemployment of coaches and missing data). The total number of potential shared-coach teams is $49\times1+12\times3+1\times6+1\times10+1\times15 = 116$. With 679 players for whom the coach is observed, the total number of potential matches with any partner (the number of edges in a fully connected set) is 230181. This implies a vanishingly small fraction of teams wherein the players share a coach of .00050. In fact, 102 of 6034 teams for whom coaches for both players is observed happen to share a coach: a fraction of .017.

Although this "back of the envelope" analysis does not control for other player characteristics, it appears that players who share a coach are more likely to play together. This includes players such as brothers Alexander and Mischa Zverev, who are both coached by Alexander Zverev Sr. In general, players sharing a coach may partner together because of the shared coach, or the same factors that led them both towards that coach may push them to play together. In the case of the former channel of causality, it remains ambiguous whether one should expect a shared-coach team to outperform a team with different coaches. There could be economies of scale that lead these teams to be more productive; or sharing a coach could help overcome search frictions, in which case one may expect a shared-coach team to be on average less productive. The following subsection approaches this question.

I.3 Could the team ability capture the coach's ability if they can better identify players that would 'fit' well, or would it end up in the idiosyncratic match quality?

Consider the average match quality ψ_{ij}, as recovered from the average residual over all observations for a team ij. Limit the sample to the subset of players for whom the coach is observed.[43] An indicator variable for teams sharing a coach captures the average

[43] Fixed effects estimates recovered from the two-way fixed effectsmodel tend not to react strongly to changes in the set of individuals included in the sample. This is because an individual's performance is strongly tied to their own average performance (averaged across teams). Correspondingly, the estimates of match effects recovered from residuals do not change greatly either. The correlation between the

Table 22: Players per coach for 2016

	(1)	(2)	(3)
	b/se	b/se	b/se
Shared coach	0.0166	0.0208	-0.0431
	(0.0429)	(0.0429)	(0.0594)
Different handedness		0.0393	0.0287
		(0.0240)	(0.0353)
Different backhand			0.0315
			(0.0317)
Constant	0.0155	0.00481	-0.00492
	(0.0105)	(0.0122)	(0.0206)
r2	0.0000123	0.000240	0.000361
N	12165	11857	5802

difference between those teams and the others. For comparison (and informed by the Interview section in the appendix to the main chapter) I generate also indicators for different handedness (right or left) and different backhand styles (one- or two-handed). Table 22 shows regressions of these indicators on the estimated residual match effects $\hat{\psi}$. None are significant at usual levels, although sharing a coach seems to predict weakly higher match quality (in the first two specifications), as does partners having different dominant hands. The sign reversal in the third specification is likely due to the change in sample size (as many entries for backhand contain missing data).

Together these results do not provide evidence of a strong role for shared coaches in determining match effects, nor handedness nor backhand styles. Neither do they rule out such channels; players may look for a variety of complementary characteristics in their partners, and trade off the observables characteristics analyzed above with unobservables. In this case the positive effects of shared coaches or different dominant hands could be underestimated due to downward bias from unobservables – players accept a partner with less complementary unobservables because of complementarity in observables.

residual match effects recovered from the full sample and those recovered from the sample restricted to teams wherein both coaches are observed is .97.

I.4 A final caveat

When interpreting the above results on shared coach matching probability and match quality, keep in mind the noisiness of the measurements: coaches are reported in 2016 while the estimation sample ranges from 2009 to 2015. The more players switch coaches, the more measurement error is introduced.

J Forfeits

When one partner withdraws from a tournament the other must also withdraw – they are admitted into and advance through the tournament as a unit. They receive the prize money corresponding to the round they attained before withdrawing, and the withdrawing player pays a small fee (while the partner does not). The opponents receive a 'walkover' and advance to the next round.

The tournament-level dataset does not report scores or indicate whether a walkover occured. A match-level dataset provided to me by the ATP contains many incidents of 'retirements' – when a player withdraws during a match due to injury or exhaustion. All of these report scores within the match up until the retirement, indicating that the match proceeded initially. No walkovers are reported. However, anecdotal evidence suggests they are not vanishingly rare, so they may be omitted from the dataset.

If strong singles players frequently withdraw from doubles competition (so as to focus on the singles category) their team skill is systematically underestimated, because their doubles tournament results do not reflect their ability to perform in this category. Likewise, the team skills of doubles specialists who do not choose to withdraw from doubles play is systematically overestimated, because they advance further in tournaments due to walkovers. This would bias the main results in favour of finding a large difference in skill factors when variations in tournament results are in fact due in part to differential effort.

Subsection 4.5 in the main text speaks to this concern. This subsection presents a recreation of the main estimation with a sample restricted to 'nonbusy' tournaments – those for which a player participates only in one category (singles or doubles). In a nonbusy tournament there is no question of withdrawing to focus on another category. The estimated skill factors are virtually identical, having correlations of .96 and .97 with the baseline skills for general and team skills respectively.

K What is being identified by the team skill estimates?

Pro tennis certainly requires a whole bundle of skills. Some of these may be more or less intensive in doubles compared to singles. Economists are interested in interpersonal skills – the ability to coordinate and cooperate – but uninterested in the particular mechanics of pro tennis. So it would be desirable to be able to isolate the skill factors of interest, but this is not feasible given the current data. For a given player there are two productivity measures – singles and doubles – and identifying more skill factors would require more within-player degrees of freedom (more on this in the following section).

To the extent that the intensity of tennis-specific athletic skills differs across the singles and doubles production functions, the skill estimates pick up differences in these skills across players. So although I would prefer a cleanly identified estimate of 'coordination skills', given the limitations of the data I estimate the broader factor of 'team skills' – the latter being defined as any ability to outperform at teamwork what would be predicted by one's solitary productivity.

I defend the contribution as follows. Previous studies have identified multiple skill factors, for multiple workers, from intermediate outputs of a team production function (see for example Arcidiacono et al. 2017); this requires assumptions on the functional form of team production. The current chapter exploits instead multiple output measurements within a player (solo work and teamwork). This necessarily introduces the problem of comparing different modes of production – so the drawback outlined above is integrally linked to the identification contribution, which should be seen as complementary to the existing literature but not without limitations.

Moreover, I argue that team skill can reasonably be interpreted as a proxy for interpersonal skills. Consider a dataset of many workers participating in solo production and team production. Each worker i has $K > 1$ skills, each given by α_{ik}, and K may be very large. Solo output is determined by the following production function:

$$y_i^1 = \sum_{k=1}^{K} \alpha_{ik} w_{sk} = \alpha_i' w_s$$

where $w_{sk} \in \mathbb{R}^+$ indicates the weight of skill k in solo production. Team output for a team ij is given by

$$\begin{aligned} y_{ij}^2 &= \sum_{k=1}^{K} \alpha_{ik} w_{dk} + \sum_{k=1}^{K} \alpha_{jk} w_{dk} \\ &= \alpha_i' w_d + \alpha_j' w_d \\ &\equiv z_i^2 + z_j^2 \end{aligned}$$

where $w_{dk} \in \mathbb{R}^+$ indicates the weight of skill k in the team production function. In a connected set of teams, both terms in the above production function can be separately identified. Then the first term can be compared to the solo output of worker i to identify team skill. Letting y^1 and z^2 represent vectors of productivity terms for I workers in population $\mathcal{I}$, the estimate of team skill is given by $\hat{\gamma} = z^2 - (y^{1'} y^1)^{-1} y^{1'} z^2 z^2$; that is, the extent to which each worker overperforms (or underperforms) at teamwork compared to what their solo productivity predicts. Consider the extreme case where weights are identical across production functions; that is, $w_s = w_d$. In this case y^1 is a perfect predictor of z^2 and team skill is equal to zero. The greater the difference in weights, the greater the estimate of team skill. Moreover, those skills k for whom $|w_{sk} - w_{dk}|$ is greatest contribute proportionally more to the estimates of team skill. These may include skills such as volleying and movement. However, I argue that $|w_{sk} - w_{dk}|$ is likely to be greatest for $k = \textit{interpersonal coordination}$, since $w_{s\textit{interpersonal coordination}} = 0$. This suggests that the estimates of team skill should be reasonable proxies for interpersonal ability.

L Data on within-match statistics

Without more within-player data, no object more disaggregated than team skill can be identified. There is some possibility of using 'micro-statistics' from within matches, as suggested above. These could potentially be used to determine the relative contribution of different skill factors between solo and team production. However, such analysis is obstructed by two factors. The first is that the micro-statistics are intermediate outcomes that contribute to productive output; they are neither skill factors nor weights, and some theory must be introduced if they are to be mapped into such objects. This would bring the analysis back into the territory of Arcidiacono et al. (2017), which makes assumptions regarding how different skills map into different intermediate outcomes. Doing so would not be in line with the identification contribution of the current chapter, but would certainly shed new light on the different underlying skills that produce the estimates of team skill. More crucially, micro-statistics are never reported for doubles matches, but only for singles. This severely limits the extent to which such analysis could be elucidating. I consider pursuit in this direction a valuable topic for future research.