Injury, Workload, and the Strategic Response of Coaches in the NBA*

Kala Krishna† Sergey Lychagin‡ Lewis S. McLean§

August 25, 2026

Abstract

Worker effort increases contemporaneous output but may result in injury or burnout. Directing effort away from the workers most exposed to injury would reduce this risk, but information on the exposure is often hidden. We ask whether managers act on such private information, studying the NBA as a data-rich workplace. Instrumenting workload with teammate absence, we find that work causally raises injury risk, which contrasts with the spurious negative correlation commonly found in sports medicine studies. We also show that coaches rest players in response to injury risk shocks. By acting on this private information, coaches can raise player productivity.

JEL Classification: D83, J28, M54, Z22

Keywords: task assignment; managerial discretion; workplace injury; workload; worker health and burnout; load management in sports

*We would like to thank Conor Ryan, Joris Pinkse, Douglas Wrenn, Paul Grieco, and participants at the 2023 Latin American Sports Economics Conference (LASEC) for insightful comments. Sergey Lychagin thanks the WIFO Institute (Vienna) for hosting him during work on this project. All remaining errors are our own.

†The Pennsylvania State University, NBER, IGC and CES-IFO, e-mail:kmk4@psu.edu

‡Central European University, Vienna, Austria, e-mail:lychagins@ceu.edu

§Analysis Group, Washington, D.C., e-mail:lewis.mclean@analysisgroup.com

1 Introduction

Some workers are valuable and hard to replace, and the work they do today can hurt their output tomorrow. A manager who could see which of them is closest to breaking would lean on that worker less; but a worker's condition is largely hidden, so whether managers hold and act on such private, time-varying information about who is at risk is far from obvious. That is the question of this paper.

Absent that information, a manager has only blunt instruments: lean hardest on his strongest workers, or tie pay to output so that workers drive themselves (Lazear, 2000). Either way, work is extracted without regard to who is at risk, and the cost is injury and stress (Artz and Heywood, 2015; Baktash et al., 2022). Information about individual condition could instead let a manager pull back on the most vulnerable worker, which returns us to whether managers actually have and use this information.

The question is hard to address, on two counts. First, the manager's information is unobserved: what looks like informed discretion can be a valuable private signal or plain bias. In hiring, Hoffman et al. (2018) find that the managers who most often override a screening test go on to hire shorter-tenured workers on average, which they read as discretion reflecting bias or mistakes. The same could be true of a coach's decision to rest a player: from the outside, an informed save and a mistake look alike. Second, and worse for measurement, when a manager reduces work on a private assessment of risk, the assignment of work becomes endogenous to the very injuries it produces: the workers he spares are the vulnerable ones, so in the data, lighter workloads and injuries occur together, and work can appear to *prevent* the harm it in fact causes. That confounding is itself part of what we study.

Professional basketball is a good place to examine this mechanism. Sport has long served as a laboratory for questions in labor economics that are hard to observe elsewhere (Kahn, 2000), and elite basketball is among the most information-rich management settings one can study: coaches, advised by dedicated medical and performance staff, hold unusually detailed, real-time information about player condition, and their fielding decisions are recorded game

by game. What the coach sees is private, but the record of his decisions is public and rich. We observe who was available for each game, how many minutes each player was used, when an injury occurred, and when an available player was held out, so a manager's decisions and their consequences can be studied together.

Our analysis centers on a single outcome: whether a player suffers a time-loss injury in a given game. We ask how the probability of that injury depends on two measures of work: his *current* workload, the minutes he plays in that game, and his *chronic* workload, the minutes he has accumulated over the preceding one to four weeks. Separating current from chronic work in this way is standard in the sports-science literature (Griffin et al., 2020).

For the current game, the effect of interest is that of the *assigned* minutes on player injury risk. Assigned minutes are the minutes the coach plans to allocate to a player before the game begins. But assigned minutes are endogenous: the coach sets them on his private assessment of who is vulnerable, giving fewer to the players he judges most at risk, so they already reflect the very risk we are trying to explain. Realized minutes, what the player actually plays, are worse still, because injury can cut an appearance short. We therefore instrument current minutes with teammate absence. Because five players must be on the court at all times, a teammate's absence mechanically pushes minutes onto those who remain, for reasons that have nothing to do with the coach's assessment of any one player. We use only this absence-driven part of the variation in minutes, which shifts assigned work while leaving the coach's private judgment out of it. We estimate the injury model using the probit control function approach.[1]

This is only half of the design: it identifies the cost of playing a player, the premise that makes load management worthwhile. On its own, though, it cannot reveal whether coaches act on private information: the naive estimate is biased in the direction we would expect if coaches were sparing the players they judge most vulnerable, but injury-truncation biases it the same way, and the two cannot be separated. The other half of the design asks what a coach reveals when he chooses to bench a player.

[1]In the appendix, we also present an analogous linear probability model estimated via 2SLS.

When a coach rests an available player, that choice can carry his private assessment of the player's condition; when a player is out for illness or a personal reason, we take the absence to be unrelated to what the coach thinks. Thus, we compare injury risk after the two types of absence: a difference between them would mean the coach was acting on something he saw and we did not. As a by-product, the same comparison settles a question the sports-medicine literature has muddled: whether rest itself raises injury, or whether coaches simply rest players who are already at risk.

Three findings follow. First, instrumenting reverses the sign of the current-work effect: realized workload is *negatively* associated with injury in naive specifications, the pattern usually reported (Caparrós et al., 2018; Gabbett and Ullah, 2012), but the effect of current work is positive and significant once it is instrumented. Our preferred estimates imply that an additional assigned minute raises a star player's injury probability by roughly a tenth of a percentage point, small per minute but cumulative over a game. Second, once current assigned work and a recent-absence proxy are included, we no longer detect a partial relationship, linear or U-shaped, between injury and chronic workload over horizons of one to four weeks. Third, and most importantly, injury risk rises after a coach rests an available player but not after an illness or personal absence, a gap that points to coaches acting on information we do not observe.

Our primary contribution is to the economics of management and private information. A large literature shows that managers matter, for output and turnover (Lazear et al., 2015), through structured practices (Bloom et al., 2013), retention and interpersonal skill (Hoffman and Tadelis, 2021), persistent traits (Adhvaryu et al., 2023), and the reallocation of tasks after transient shocks (Adhvaryu et al., 2022). We add evidence that managers act on private, time-varying information about a worker's *state*, the kind of signal no personnel dataset records, and that failing to account for how work is assigned inverts the apparent relationship between work and injury. The closest antecedent, Adhvaryu et al. (2022), studies managers reacting to a shock the analyst also sees, ambient pollution, by reshuffling workers across tasks to protect current output; their managers respond to output they watch fall,

bundles piling up on the line, not to a forecast of which worker is about to break down. Our subject is that forecast itself: private information about one worker's vulnerability, used to protect tomorrow's availability by reducing today's output. In that respect the paper is related to work that uses administrative data to recover how experts' private judgment shapes the decisions they make (Currie and MacLeod, 2017; Abaluck et al., 2016).

The mechanism is not special to basketball. The value of acting on private information about worker state is greatest where the best workers are highly skilled and hard to replace, as in our setting, but it is present wherever heavy workloads risk costly mistakes, absenteeism, or turnover. Doctors and others working in high-pressure environments are at risk of burnout, which is both prevalent and costly: Prasad et al. (2021) report burnout among 49 percent of US healthcare workers surveyed during the COVID-19 pandemic, and Nekoei et al. (2025), using Swedish data, estimate that burnout reduces economy-wide labor earnings, directly and indirectly, by 3.6 percent. Our results suggest that a manager can manage workload to reduce it. The value of such information will only grow as more workplaces monitor their workers as thoroughly as the NBA already does (Milanez et al., 2025). Because managers here act on private, hard-to-verify signals, the paper also connects to work on incentives and subjective information inside firms (Prendergast, 1999) and to the personnel economics of task assignment (Burgess et al., 2010; Amodio and Martinez-Carrasco, 2018).

The paper also relates to quasi-experimental evidence that working time affects health and safety. Using statutory reductions in the workweek, Lee and Lee (2016) find that fewer weekly hours lower reported injuries per hour worked, and Cygan-Rehm and Wunder (2018) find evidence that longer hours worsen self-assessed health; Berniell and Bietenbeck (2020) find that a shorter workweek reduces smoking. That evidence is complementary: those studies identify the effect of institutionally mandated hours, whereas we study a manager's discretionary, information-driven allocation of work.

Finally, because our outcome is injury, the results speak directly to sports-medicine research on workload. That literature attributes the raw associations between workload, rest, and injury to biomechanics. But its estimates are subject to the two problems identified

above: they rest on realized exposure, which is simultaneous with injury, and on workloads that coaches assign by selecting on hidden risk. A systematic review (Griffin et al., 2020) treats the acute-to-chronic workload ratio as associated with non-contact injury, and a reanalysis (Impellizzeri et al., 2021) shows the ratio's construction can manufacture such associations on its own. The literature reports two such patterns: current work appears to reduce injury risk, a result the field itself finds puzzling, and chronic workload has a U-shaped relationship with injury, the field's consensus. Our results show that both can arise from how work is truncated by injury and assigned by coaches selecting on hidden risk, rather than from biomechanics alone.

The remainder of the paper proceeds as follows. Section 2 describes how coaches monitor player condition and reviews the specialized sports-medicine literature on workload and injury, along with the econometric problems it leaves unaddressed. Section 3 describes the data. Section 4 develops the probit control-function strategy and its assumptions. Section 5 presents our three main results, and Section 6 concludes.

2 Background

2.1 Monitoring player condition in the NBA

A modern NBA coach does not judge player condition alone. Every team surrounds its coaching staff with a dedicated medical and performance department: team physicians, athletic trainers, physical therapists, strength and conditioning coaches, and dietitians, often ten or more people overall (National Basketball Association and National Basketball Players Association, 2023; Trudell, 2018). These staff are in near-daily contact with the players, who practice, travel, and recover with the team across an eighty-two-game season. Each morning, the medical staff reviews every player's status, what he did the day before and what is planned for the day ahead, and relays that assessment to the coaches, so that decisions about who plays and for how long are made against a current read on each player's body (Trudell, 2018).

What that read draws on goes well beyond anything visible in a box score. Players report daily on soreness, fatigue, and sleep, and staff watch them in practice and warmups. Teams supplement these subjective signals with objective monitoring. During practice, athletes wear GPS and accelerometer units that record distance, acceleration, and the mechanical load of each movement. Wearable monitors track heart-rate variability, sleep, and recovery, and force plates and strength testing gauge readiness and fatigue. In games, the league's optical tracking system, cameras mounted in the catwalk of every arena, records the position of every player many times per second. A coach can therefore see when a player's movement is dropping off, when his recovery is incomplete, or when accumulated load is running ahead of what he can absorb.

This information routinely shapes what coaches do, in ways anyone who follows the league recognizes. Teams place returning players on minutes restrictions and ramp their workloads up gradually; they rest healthy stars on the second night of a back-to-back or through congested stretches of the schedule, a practice now common enough to have acquired its own label, load management. Such choices hinge on judgments about condition and risk that the coaching and medical staff can make but an outside observer cannot. In short, coaches hold detailed, continuously updated, and largely private information about which of their players is most at risk of injury, and they act on it.

Yet whether acting on that information helps is disputed. The NBA maintains, on ten years of league-wide injury data, that resting players does not lower their risk of injury (Powell, 2024), and since 2023–24 its Player Participation Policy has fined teams for resting healthy stars. Whether coaches read risk correctly, and with what consequences for injury, is the question this paper takes up.

2.2 Workload and injury in sports medicine

Sports medicine research on injury prevention and recovery focuses heavily on workload, which it treats as a key modifiable determinant of injury risk (Griffin et al., 2020). The aim of this literature is to help coaches, physicians, and athletes structure training programs

and playing schedules so as to reduce injury risk while sustaining top performance. Because workload is one of the few determinants of injury that a coach can set directly, it occupies a central place in the field.

The relationship between workload and injury depends on the horizon over which work is measured. Measured instantaneously, work is the application of force to tissue, which simple biomechanics identifies as a direct cause of acute injury. Measured over longer horizons, the interplay among stress, recovery, and adaptation becomes central: the principle of progressive overload holds that gradually increasing the volume of work builds endurance and strengthens musculoskeletal tissue (Kraemer et al., 2002; Soligard et al., 2016), while adaptation requires adequate rest for the body to recover (Soligard et al., 2016). Together, these forces imply that the workload-injury relationship need not be monotonic at longer horizons.

Researchers have studied many measures of workload in search of strategies to reduce injury risk, and the resulting evidence is varied and at times perplexing. Studies that model injury as a function of current realized work are notable for reporting *negative* relationships, in which athletes who perform less work are observed to be injured more often rather than less (in particular, see Gabbett and Ullah, 2012; Caparrós et al., 2018). This pattern is conspicuous, because the marginal effect of work on tissue should, by simple biomechanics, raise the risk of injury rather than lower it.

Chronic workload, defined as total work accumulated over one to four weeks, has been studied just as thoroughly, again with mixed results: reported associations with injury include null findings (Wong et al., 2022), positive associations (Lewis, 2018; Caparrós et al., 2016; Orchard et al., 2009; Menon et al., 2024), and negative ones (Hulin et al., 2014), while Cross et al. (2016) report a U-shaped pattern for four-week loads and Bache-Mathiesen et al. (2021, 2022) argue that the relationship is non-linear and should be modeled as such. Despite this heterogeneity, the prevailing view holds that injury risk is U-shaped in chronic workload, falling at first as baseline fitness improves and rising once workloads exceed the level that allows adequate recovery (Griffin et al., 2020). The same view is applied to the acute-to-chronic workload ratio, short-term work divided by chronic work (Griffin et al., 2020). The

ratio has itself drawn scrutiny from within the field, as Impellizzeri et al. (2021) show that expressing injury risk as a function of such a ratio can manufacture associations that reflect the arithmetic of the ratio rather than any workload mechanism.

These physiological mechanisms are real, and the measurements behind them are careful. What the estimates lack is a research design that confronts the endogeneity of workload, of the kind set out in Section 1. For current work, endogeneity takes two forms. The first is simultaneity: an injury typically cuts short a player's participation, so a low realized workload can be a consequence of injury as much as its determinant. The second is endogenous assignment: coaches may give fewer minutes to the players they judge most vulnerable. The few studies that acknowledge the simultaneity address it only partially, relating workload to injury in the following week rather than the current one (Hulin et al., 2014), or fitting survival models on data partitioned by the amount of work performed, such as the distance run at a given pace in Gabbett and Ullah (2012) or accelerometer-derived player load in Wong et al. (2022). Because the partitioning variable is itself realized work, and so endogenous with respect to injury, these designs do not recover the effect of interest.

Chronic workload, being predetermined, escapes the simultaneity issue but faces two omitted-variable problems that the literature leaves unaddressed. First, chronic workload is highly correlated with current assigned work, which this literature does not observe; the studies that do control for current work control for work performed, which is itself truncated by injury. Second, it is confounded by persistent injury-risk shocks: recent injuries mechanically depress chronic workload, and coaches, drawing on the kind of private information described above, assign lighter loads to the players they judge most at risk. Both channels push chronic workload and injury together in ways unrelated to any biomechanical effect.

Taken together, this literature estimates the relationship between work and injury without a design that separates the effect of assigned work from the confounding introduced by injury and by coaches' responses to it; in the words of one review, "no studies have even tried to estimate causal effects properly" (Impellizzeri et al., 2020).

3 Data Overview

For this study, we constructed a comprehensive player injury and workload panel dataset encompassing all NBA games from the start of the 201718 season to the end of the 202223 season. This dataset integrates detailed information on player workload, demographics, and injuries for $1,126$ professional basketball players over a total of $7,581$ official games. It contains a total of $251,037$ player-game observations, consisting of $160,615$ instances in which a player was fielded and $90,422$ instances in which they were benched (i.e., assigned zero minutes for that game). All data were obtained from public sources, the four primary sources being NBA.com, ProSportsTransactions.com, FoxSports.com, and ESPN.com.

For data detailing player workload, we scraped box-score data from NBA.com, which detail appearances and minutes played for each player in every official game during the sample period. We define current realized work for player i in game g as total minutes played (MP_{ig}) during each game. Using these data, we also construct measures of chronic workload as weekly-average hours played at horizons of one, two, three, and four weeks. We abbreviate these chronic workload measures by $WHP_{ig}(X)$, which denotes "weekly-average hours played at the X-week horizon" for player i as of game g. To define $WHP_{ig}(X)$ precisely, let $\mathcal{M}_{ig}(X)$ denote the set of official games played by player i's team over the last X weeks as of, but excluding, game g.[2] Then, we construct $WHP_{ig}(X)$ as

$$\text{WHP}_{ig}(X) \equiv \sum_{s \in \mathcal{M}_{ig}(X)} \frac{\text{MP}_{is}}{60 \cdot X}, \tag{1}$$

where the division by $60 \cdot X$ gives this metric its "hours per week" interpretation. Note that under this definition, $WHP_{ig}(X)$ is necessarily weakly monotonically increasing from each player's first appearance of the season until X weeks have passed. It is thus also weakly monotonically increasing for the first X weeks of the season. This is appropriate given that workload consists entirely of time spent participating in official games, which implies that

[2]In other words, $\mathcal{M}_{ig}(X)$ is the set of all official games in which player i could possibly have participated over the past X weeks as of, but excluding, game g.

	Minimum	25th Percentile	Median	Mean	75th Percentile	Maximum
MP_{ig}	0	15.20	23.70	22.80	31.37	64.97
$WHP(1)$	0	0.51	1.03	1.02	1.51	3.00
$WHP(2)$	0	0.51	1.00	0.99	1.46	2.80
$WHP(3)$	0	0.48	0.98	0.96	1.42	2.79
$WHP(4)$	0	0.44	0.94	0.93	1.39	2.61

Table 1: Descriptive statistics for different measures of current and chronic workload.

no work is performed prior to the start of a season. Descriptive statistics for MP_{ig} and $WHP_{ig}(X)$ are outlined in Table 1.

Player i's game g injury outcome is a binary variable denoted by $Injury_{ig}$. Injury data were obtained via web scraping from two primary sources: ProSportsTransactions.com (PST) and FoxSports.com (FS). PST provides a database documenting NBA player transactions, including injuries, that extends back to the start of professional basketball (Marousek, 2005). It is, to the best of our knowledge, the most complete public online database of NBA player injuries in existence, has been used to study injuries among NBA players on several occasions [see for instance Lewis (2018) and Bullock et al. (2022), among others], and includes both injury-onset and return-to-active-roster dates.

To ensure its accuracy, we cross-checked the timing and duration of injuries recorded in PST against the NBA.com box-score data, which provides player status on a game-by-game basis (e.g., active, inactive, or out due to injury). This process allowed us to validate and, when necessary, adjust reported durations where discrepancies were found. The FS data, which are less complete than the PST data because they record only injury occurrences and games missed due to injury, were then incorporated to identify any additional injuries not captured in PST. Since FS does not include injury durations, we inferred them by analyzing gaps in player availability using the box-score data. These datasets were merged with the workload data by date. A detailed description of the cleaning and matching process is available in Appendix A.

Consolidating the injury data yielded a total of $6,570$ unique time-loss injury events, i.e., injuries that caused a player to miss one or more games. Of these, $6,161$ were found in

the PST data and 4,958 were found in the FS data, with an overlap of 4,549 injury events between PST and FS. Of these, 5,258 were matched by date to a specific game, yielding the binary variable $Injury_{ig}$.[3] Dividing $Injury_{ig}$ by the total number of player appearances yields an injury rate of 32.74 instances per 1,000 appearances, or a one-in-thirty chance of a time-loss injury per appearance.

Injury duration—defined as the time between an injury event and the player's return to the active roster—is heavily right-skewed in our sample, with a median duration of six days and a mean duration of 18.8 days.[4] Alternatively, injury duration can also be measured in terms of the number of missed games caused by the injury. By this metric, the median injury duration is two games and the mean injury duration is 6.2 games.

In total, we record 90,422 instances in which a player was not fielded during one of their team's games. We categorize these instances of absenteeism into four broad groups: *Coach*, *Injury*, *Illness*, and *Other*. *Injury* and *Coach* are the two largest groups and account for 40,444 and 44,734 missed games, respectively. *Injury* specifically refers here to all manner of physical injuries that resulted from competitive play, practice, or training.[5] *Illness*, which includes all manner of viral and bacterial illnesses and all medical conditions that were not deemed to be play-related (e.g., wisdom-tooth removal), is the third-largest category and accounts for a total of 4,237 missed games. Of these, a total of 2,413 instances of absenteeism were attributed to COVID-19 health-and-safety protocols, under which players were ordered to quarantine after a positive COVID-19 test. The smallest category, *Other*, accounts for 1,007 missed games and covers personal and technical reasons such as family emergencies and suspensions for player misconduct. All instances for which there was no information indicating that they belonged to one of the above categories, or in which the player was

[3]The 903 observations that could not be matched had start dates falling after a game in which the player was not fielded, implying that the player was most likely injured during practice or while participating in a lower-league game, as is common among less-skilled players, rookies, and players undergoing long-term rehabilitation. More information about the cleaning and matching process is available in Appendix A.

[4]As we are working with time-loss injuries, the minimum injury duration is one day by definition. The maximum injury duration recorded in our sample was 337 days.

[5]Note that this is more expansive than the set of injuries with onset captured by $Injury_{ig}$, where the latter includes only those injuries that could be matched to a specific game.

designated as out for rest (1,125 missed games), are allocated to *Coach*; a discretionary decision by the coach not to field a player would fall under this category.

In addition to distinguishing instances in which coaches chose not to field available players, we also use these player-absenteeism categories to construct an "exogenous teammate absenteeism" variable to serve as an instrument for current work in our reduced-form analysis. This variable, denoted Z, is constructed as follows: Let $Absent_{jg}$ be an indicator variable that takes on a value of one if player j was out due to *Injury*, *Illness*, or *Other* reasons during game g, and let $\mathcal{J}_g$ be the set of players on player i's team. The instrument Z is constructed as

$$Z_{ig} \equiv \sum_{j \in \mathcal{J}_g} \text{Absent}_{jg} \cdot \mathbb{1}(j \neq i).$$

In other words, Z_{ig} is the number of player i's teammates who were not available to participate in game g. Note that players who were available during the game but were not fielded by the coach (i.e., instances in the *Coach* group) are not counted as absent and thus do not contribute to Z_{ig}. This is because coaches endogenously determine whether or not to field *available* players.

Lastly, we collected season-level salary data from ESPN.com, which we use to construct a season-level categorical salary-decile variable, denoted $Decile_{ig}$.[6] For some players, salary information is unavailable in our ESPN data—typically because they were on temporary two-way or 10-day contracts—and we treat them as an 11th "decile".[7] Unsurprisingly, $Decile_{ig}$, which is a proxy for player skill, is highly predictive of both player workload and the incidence of injury. In our analysis, we use $Decile_{ig}$ in lieu of player fixed effects to control for correlation between assigned work and compositional differences in baseline risk across the skill distribution, which is necessary given the nonlinear model used in the analysis.

[6] $Decile_{ig}$ varies across seasons and is fixed within seasons. We include the g subscript to denote variation across seasons without introducing an additional subscript.

[7] A few established players lack a matched salary in some season—for instance, LeBron James in the 2017–18 season. For each such player we assign the salary decile averaged from their nearest matched seasons, so that they no longer fall in the 11th-decile bin.

4 Identification, Estimation and Inference

Throughout our analysis, we use the same basic probit control function identification strategy proposed by Rivers and Vuong (1988) to identify the causal effect of current workload MP_{ig} on injury. As discussed later, we repeat our main control-function analysis using the linear probability model (LPM) and two-stage least squares. The LPM results are largely the same as those we obtain using the control-function approach.

First, we outline the basic identification strategy. Define $x_{ig} = f(w_{ig}, MP_{ig})$ as a vector consisting of control variables w_{ig} and endogenous current realized work MP_{ig} transformed by the function $f(\cdot)$. As a special case, $f(w_{ig}, MP_{ig})$ simply returns the vector of covariates (w_{ig}, MP_{ig}), but $f(w_{ig}, MP_{ig})$ could also return interactions between w_{ig} and MP_{ig} or nonlinear transformations of MP_{ig}. The model of interest is specified as

$$\text{Injury}_{ig} = \mathbb{1}(x_{ig}\beta + \epsilon_{ig} > 0) \tag{2}$$

where $\mathbb{1}(\cdot)$ is the indicator function. Assuming the injury shock ϵ_{ig} is standard normal makes (2) a probit model, but the parameters are not yet identified since x_{ig} and ϵ_{ig} are not independent.

Next, define the column vector $q_{ig} = (w'_{ig}, Z_{ig})'$, where Z_{ig} denotes the teammate-absenteeism instrument and is excluded from (2). The endogenous variable MP_{ig} can then be expressed as

$$\text{MP}_{ig} = q'_{ig}\pi + e_{ig} \tag{3}$$

for a conformable vector of first-stage parameters π. Assuming (ϵ_{ig}, e_{ig}) is jointly normal and independent of q_{ig}, i.e., $(\epsilon_{ig}, e_{ig}) \perp\!\!\!\perp q_{ig}$, it follows that

$$\Pr(\text{Injury}_{ig} = 1 | x_{ig}, e_{ig}) = \Phi(x_{ig}\widetilde{\beta} + e_{ig}\widetilde{\gamma}) \tag{4}$$

where $\epsilon_{ig} - e_{ig}\gamma \mid e_{ig} \sim N(0,\ \sigma^2)$ and $\widetilde{\beta} \equiv \beta/\sigma$. This is the probit control function, or

probit-CF.

A few comments are in order. First, as with 2SLS, second-stage identification relies on instrument relevance and (a strengthened version of) instrument validity. Instrument relevance is all but assured by the constraints inherent to the environment and simple arithmetic: During the 48 minutes of regulation in each game g, the coach must field no fewer than five players at all times.[8] Thus, as additional players become absent, the remaining players will tend to be assigned additional minutes on-court.

Given a probit second stage, the joint independence assumption $(\epsilon_{ig}, e_{ig}) \perp\!\!\!\perp q_{ig}$ serves as the instrument validity assumption. Ignoring the other controls in q_{ig}, this condition implies that teammate absenteeism Z_{ig} is independent of the distribution of player injury except through its effect on the conditional expectation of MP_{ig}. It is similar to but significantly stronger than the usual orthogonality assumption required for identification via 2SLS, and is required for identification whenever $f(w_{ig}, MP_{ig}) \neq (w_{ig}, MP_{ig})$. Although some results below are consistent under weaker assumptions, we maintain the independence assumption throughout our analysis for the sake of internal consistency and brevity. Appendices B.2 and B.3 provide linear probability model (LPM) counterparts to each probit-CF estimate; these counterparts are identified under a weaker assumption whenever $f(w_{ig}, MP_{ig}) = (w_{ig}, MP_{ig})$ and yield similar results.

Second, note that the coefficients $\widetilde{\beta}$ in (4) are, in general, not equal to β in equation (2) and thus are not directly comparable between model specifications. While we can identify β and σ separately,[9] inference is technically complicated; details on inference can be found in Appendix C. In the body of the paper, we report $\widetilde{\beta}$ as well as $1/\sigma$ so that the estimates can be compared across different specifications. We also report marginal effects to convey the magnitude of the effects.

[8]During regulation, no player can play more than 48 minutes. If, however, the game is tied at the end of regulation, it goes into a five-minute overtime period. This process is repeated until the point spread is nonzero at the end of an overtime period. Thus, if Overtimes_g is the number of overtimes in game g, then $48 + 5 \times \text{Overtimes}_g$ is the maximum number of minutes any one player can spend on-court. The current record for the number of overtimes in a single NBA game is six, which was set on January 6, 1951, in a game between the now-defunct Indianapolis Olympians and the Rochester Royals.

[9]This is shown for univariate e by Wooldridge (2015), among others.

Finally, basic physiology—and indeed, some of our findings—suggests that the probability of injury is persistent across time within individuals, which is inconsistent with the usual MLE assumption of independence across observations. Moreover, our findings indicate that this persistence is time-varying within players. We deal with this by treating our estimator as a quasi-MLE and clustering standard errors at the player-season level.[10] For every model reported in this section (and in Appendices B.2 and B.3), standard errors are clustered at the player-season level.

5 Results

5.1 Injury and Current Work

The sports medicine literature finds that current work has a negative association with player injury; see Caparrós et al. (2018) on men's professional basketball and Gabbett and Ullah (2012) on rugby. This is counterintuitive. A positive relationship is expected since an on-court injury can occur only when a player plays, and because work causes injury. However, a negative estimated injury-workload relationship could easily result from coaches assigning less work to at-risk players (selection bias) or because injury cuts short a player's time on the court (simultaneity). In this section, we demonstrate that this negative spurious correlation exists in the data and vanishes after we instrument current work.

The primary model studied in this section (i.e., the specific parameterization of model (2)) is a simple univariate probit

$$\text{Injury}_{ig} = \mathbb{1}(\beta_0 + \beta_1 \text{MP}_{ig} + \epsilon_{ig}), \tag{5}$$

[10]This approach is supported by Hansen and Lee (2019), who show that clustered standard errors yield valid statistical tests so long as the quasi-MLE is consistent and obeys certain regularity conditions. Building on their clustered-sample asymptotics, which allow for clusters of unequal size, we obtain an analytic form for clustered quasi-MLE standard errors with generated regressors, which is available in Appendix C. This result, which is a straightforward modification of a similar one given by Wooldridge (2010) for the probit-CF without clustering [originally due to Murphy and Topel (2002)], allows for arbitrary within-group dependence and, unlike a block bootstrap, is computationally efficient.

wherein MP_{ig} is again current realized work measured in minutes played for player i in game g, and β_1 is the coefficient of interest. As a secondary model of interest, we also estimate a modified version of (5) wherein MP_{ig} and its square are both included in the model, i.e.,

$$\text{Injury}_{ig} = \mathbb{1}(\beta_0 + \beta_1 \text{MP}_{ig} + \beta_2 \text{MP}_{ig}^2 + \epsilon_{ig}). \tag{6}$$

We study this augmented specification for two reasons. First, the hypothesis $\beta_2 > 0$ is in itself of interest to the sports medicine literature, since some researchers have argued that injury risk might respond disproportionately to "workload spikes" even in a single game or competitive event [see for instance Hulin et al. (2014)]. Second, it is important to check for nonlinearity with respect to MP_{ig} to ensure that neglected nonlinearity in current work is not spuriously attributed to chronic workload in later model estimates.

Estimates of models (5) and (6) are displayed in Table 2. The first four specifications are divided into two sets, with the first two consisting of estimates of the probit model (5) and the second two consisting of probit-CF estimates of the same model. We estimate model (5) both with and without salary-decile and season fixed effects in each case. Salary-decile fixed effects are important because salary decile is correlated with both workload and injury. Season fixed effects help address the idiosyncrasies of the COVID seasons. The last two specifications are both probit-CF estimates of the quadratic workload model (6), also obtained with and without salary-decile and season fixed effects. The estimated salary-decile and season fixed effects are omitted from Table 2 for brevity and are instead shown in Figure 1 with 95% confidence intervals. Standard errors are heteroskedasticity-robust, clustered at the player-season level, and adjusted for the generated regressor $\hat{e}_{ig}$. In addition to the coefficient estimates, the panel beneath them reports the implied marginal effect of MP_{ig} on the probability of injury, in percentage points.

Turning first to the probit estimates of model (5) without the control $\hat{e}_{ig}$, both specifications yield negative and statistically significant coefficient estimates. That the sign of the coefficient is unchanged is not surprising. As expected, adding fixed effects to the model does not mitigate the endogeneity in MP_{ig}, since unobservable injury risk shocks are not fixed

	Probit	Probit	Probit-CF	Probit-CF	Probit-CF	Probit-CF
	(1)	(2)	(3)	(4)	(5)	(6)
MP	-0.006^{***}	-0.016^{***}	0.015^{*}	0.014^{***}	0.008	0.011^{**}
	(0.001)	(0.001)	(0.009)	(0.004)	(0.009)	(0.005)
MP^2					0.0002^{***}	0.0001
					(0.0001)	(0.0001)
$\hat{e}$			-0.021^{**}	-0.031^{***}	-0.022^{**}	-0.031^{***}
			(0.009)	(0.004)	(0.009)	(0.004)
Constant	-1.705^{***}	-1.701^{***}	-2.189^{***}	-2.078^{***}	-2.150^{***}	-2.065^{***}
	(0.019)	(0.046)	(0.206)	(0.074)	(0.206)	(0.075)
Marginal effects on Pr(injury), percentage points:						
MP	-0.045^{***}	-0.211^{***}	0.112	0.113^{***}	0.123^{*}	0.115^{***}
Decile & Season FE		✓		✓		✓
$1/\sigma$			1.025	1.037	1.028	1.037
IV F-Stat			69.24	603.38	69.24	603.38
Observations	160,615	160,615	160,615	160,615	160,615	160,615

Note: Marginal effects on the probability of injury, in percentage points, evaluated at a 10th-decile player in the base season with continuous regressors at their medians. *p<0.1; **p<0.05; ***p<0.01.

Table 2: Probit estimates of risk of injury with respect to current work.

and since fixed effects cannot correct for simultaneity between injury and current work.

Turning now to the probit-CF specifications (3) through (6), several features of these estimates are noteworthy. First and most important is that every probit-CF specification yields a positive coefficient estimate for MP_{ig}. Second, when MP_{ig} enters the model linearly, adding salary-decile and season fixed effects causes the standard error on MP_{ig} to fall by nearly half and substantially strengthens the model's first stage (as measured by the first-stage F-statistic), moving the coefficient from marginal significance at the ten percent level to significance at the one percent level.

Third, specification (4) implies a small risk of injury for any one individual in any given game. For scale, note that when MP_{ig} is at its median of 23.7 minutes (see Table 1), specification (4) implies a marginal effect of 0.0011 with respect to MP_{ig} for a 10th-decile player in the base season. Alternatively, for the same subpopulation of players, the model implies that a 48-minute assigned workload—i.e., an event in which a coach planned to field a player for the entirety of regulation time—results in an injury probability of 0.075.

Fourth, when MP_{ig}^2 is added to the model without fixed effects, its estimated coefficient is positive and significant at the one percent level, seemingly supporting the workload-spike hypothesis. But adding fixed effects to the model erodes this evidence, causing a more than threefold reduction in the coefficient estimate for MP_{ig}^2, rendering it insignificant, and again producing a statistically significant coefficient for MP_{ig}.

In addition to the coefficients on MP_{ig}, it is worth briefly commenting on the model fixed effects presented in Figure 1. First, the salary-decile fixed effects show that when model (5) is estimated without the control $\hat{e}_{ig}$, it fits the data by assigning a higher baseline probability of injury to higher-decile players; in many instances, the difference in the baseline probability of injury between deciles is significantly different from zero. By contrast, adding the control substantially compresses the salary-decile fixed effects, leaving only marginal differences at conventional significance levels.[11] This change indicates that the probit-CF

[11]Note that some of these salary-decile fixed effects are significantly different from zero, indicating that players without matched salary data, who form the base category, have a higher baseline probability of injury than low- to mid-decile players.

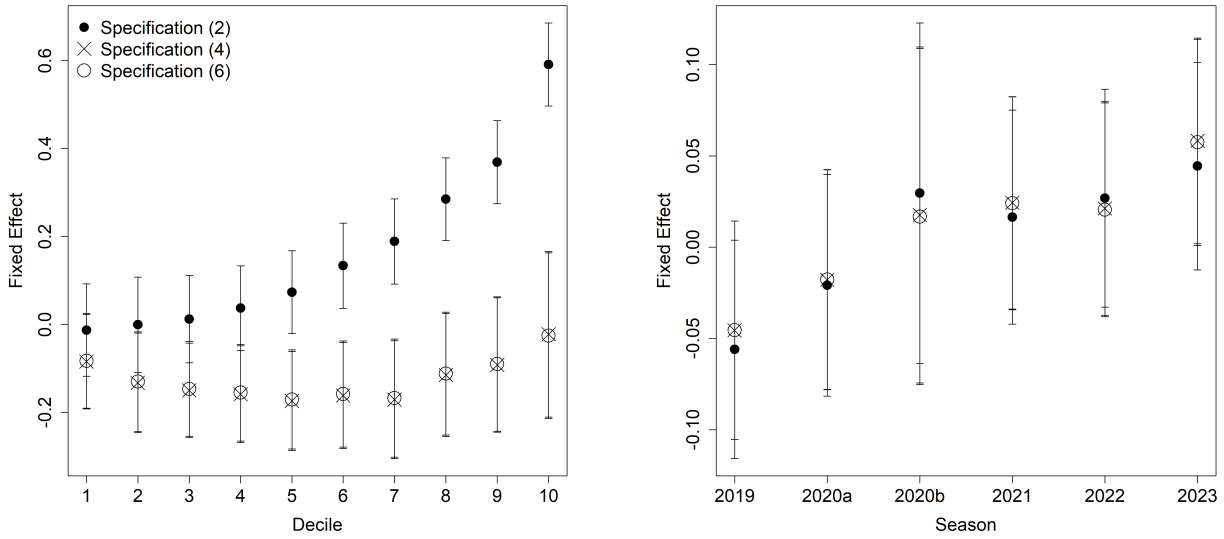

Figure 1: Salary-decile and season fixed effects with 95% confidence intervals from specifications in Table 2.

estimates attribute differences in injury rates across deciles to differences in assigned work. This is encouraging because, if the probit-CF estimates instead produced salary-decile fixed-effect estimates that increased with player decile, it might indicate underreporting of injuries among lower-decile players.

Regarding season fixed effects, two things are worth noting. First, these fixed effects are very similar across all specifications, so season fixed effects are likely not an important component of the model insofar as they control for any spurious correlation between current work and injury. Second, the upward trend in these coefficient estimates indicates that the baseline probability of injury is trending upward over time. However, the differences between these coefficients do not appear to be statistically significant in most cases, so this trend could plausibly be a result of sampling variation.

The estimates of models (5) and (6) presented in Table 2 illustrate the importance of correcting for endogeneity when modeling the marginal risk of injury with respect to current work and highlight the fundamental flaw of the sports medicine literature. All papers in this literature that we are aware of estimate the injury-workload relationship without instrument-

ing for current work, and these studies tend to find results similar to those of specifications (1) and (2) in Table 2, wherein current work is negatively associated with injury.[12] Note also that the evidence we present here provides no support for the workload-spike hypothesis, since the coefficient on MP^2 is not statistically significant in (6), our preferred specification.

5.2 Injury and Chronic Workload

Recall from Section 2 that there are two potential sources of omitted variable bias in estimating how chronic workload affects injury risk that have not been addressed in the literature. The first of these is current assigned work, which is an omitted variable in every relevant sports medicine study we are aware of. We capture current assigned work here by including MP_{ig} in the model and instrumenting it with Z_{ig}.

The second source of omitted variable bias stems from persistent injury risk shocks. One well-documented source of injury risk variation is past injuries, which have been shown to heighten future injury risk [see for instance Fulton et al. (2014) or Toohey et al. (2017)]. Since recent injuries also directly reduce chronic workloads, they will tend to generate a spurious negative correlation between low chronic workloads and injury risk. Hence, at a minimum, one should control for recent injuries when regressing player injury on chronic workload.

Another point of concern here is that coaches might have private information about player injury risk and may act on this information to help prevent injuries. Under these circumstances, some variation in chronic workload is a direct result of coaches reacting to variation in player injury risk, which we refer to as "stress shocks."[13] Supposing that stress shocks cause a persistent increase in player injury risk that decays back to baseline over time, coaches could induce a spurious negative correlation between injury risk and chronic

[12] As shown in Appendix B.2, switching from the probit control-function approach to a linear probability model estimated using 2SLS (see Table 11) results in qualitatively similar estimates. Adding player fixed effects in Table 21 also leaves the results largely unchanged, which is reassuring.

[13] One way to view this is that players could be in two unobserved states, high risk and low risk, and a stress shock puts the player into a high risk state, where an injury is more likely to occur.

workload by assigning low workloads to players when they are at their highest risk of injury and gradually increasing players' workloads as their injury risks decrease.

To control for both of these sources of persistent injury risk shocks, we construct a binary variable $Rested_{ig}$, which takes on a value of one if a player's prior-game absence was due to injury or was coach-allocated. This yields a parsimonious control for potentially heightened injury risk immediately after a player returns to the court. It is quite limited, as it neglects any persistence in player injury risk that extends beyond the first game after a return from absence, but is nonetheless effective, as demonstrated below. One should also note that $Rested_{ig}$ excludes instances in which players were absent from a game due to illness or personal/technical reasons; i.e., instances in which a player's absence was exogenous with respect to recent injuries and stress shocks.

Our examination of chronic workload thus commences with estimates of the probit model

$$\text{Injury}_{ig} = \mathbb{1}(\beta_0 + \beta_1 \text{WHP}_{ig}(X) + \beta_2 \text{Rested}_{ig} + \beta_3 \text{MP}_{ig} + \epsilon_{ig}), \tag{7}$$

wherein chronic workload $WHP_{ig}(X)$ [as defined in Equation (1)] enters the model's linear index as a single linear function.[14]

Table 3 shows estimates of the marginal risk of injury attributed to a player's one-week chronic workload $WHP_{ig}(1)$ for numerous specifications of model (7). All specifications include salary-decile and season fixed effects, with standard errors clustered at the player-season level.[15] The panel beneath the coefficient estimates reports the implied marginal

[14]For brevity, we focus this section on results obtained for $X = 1$; i.e., chronic workload measured at a one-week horizon. In Appendix B, we show results for the fully parameterized and instrumented (i.e., most robust) specifications of model (7) for $X = 2, 3, 4$.

[15]Not controlling for player fixed effects could be viewed as a limitation: the prior literature has argued that it is deviations from a player's workload—whether high or low—that put them at heightened risk [see for instance Gabbett (2016)]. As a robustness check, we therefore estimate LPM versions of models (7) and (8) that include player, salary-decile, and season fixed effects; these estimates are included in Appendix B.3. They are generally qualitatively similar to those obtained when player fixed effects are excluded. The only exceptions to this characterization are the estimates for chronic workload at the four-week horizon in Table 28 and for its square in Table 29. These coefficients are positive and significant at the five percent level, but they are economically small: at the column-(4) estimates of Table 28, five extra full games over four weeks are associated with the same increase in injury risk as about 4.5 additional minutes of play in a single game.

effects on the probability of injury, in percentage points. Specification (1), which is included as a naive baseline, yields a coefficient estimate for $WHP_{ig}(1)$ that is positive and very precisely estimated. However, it is also small: For scale, a move from the 25th to the 75th percentile of the distribution of $WHP_{ig}(1)$ (i.e., a one-hour increase in chronic workload) is associated with a 0.00423 increase in the probability of injury for 10th-decile players in the base season (i.e., the players with the highest estimated baseline risk of injury). By contrast, specification (4) in Table 2 implies an increase of 0.0182 in the probability of injury when MP_{ig} goes from the 25th to the 75th percentile for the same subpopulation (an increase of just over 15 minutes in current assigned work).

Next, we add $Rested_{ig}$ to the model to account for spurious correlation between $Injury_{ig}$ and $WHP_{ig}(1)$ resulting from selection on persistent injury risk shocks. The result, which is seen under specification (2), is a more than twofold increase in the coefficient on $WHP_{ig}(1)$, implying a substantial negative correlation between $WHP_{ig}(1)$ and instances in which a player missed the previous game due to injury or was benched by the coach. The coefficient on $Rested_{ig}$ is also large, positive, and statistically significant, capturing substantial variation in player injury. For scale, these estimates imply that $Rested_{ig}$ is associated with a 0.0268 difference in the probability of injury when $WHP_{ig}(1)$ is set to its median for 10th-decile players in the base season. This exercise reveals that there is substantial omitted variable bias in chronic workload estimates presented in the literature.

Turning next to specification (3), we add MP_{ig} to the model and instrument it with Z_{ig}, which amounts to controlling for current assigned work. Doing so yields a negative coefficient on $WHP_{ig}(1)$ (i.e., a change in sign compared with the first two specifications), and a positive coefficient on MP_{ig}. This result is again easily explained in terms of the mechanisms outlined previously: By adding MP_{ig} to the model and appropriately instrumenting it, we have controlled for current assigned work, and the corresponding reduction in the coefficient for $WHP_{ig}(1)$ is consistent with the hypothesized positive omitted variable bias.[16] Moreover,

[16] $WHP_{ig}(1)$ and MP_{ig} are positively correlated, and MP_{ig} is positively related to player injury, implying a positive bias on the coefficient estimate for $WHP_{ig}(1)$ when MP_{ig} is omitted, as observed.

	(1)	(2)	(3)	(4)	(5)
WHP(1)	0.041^{***}	0.092^{***}	-0.117^{**}	-0.051	-0.018
	(0.012)	(0.013)	(0.056)	(0.052)	(0.075)
$\text{WHP}^2(1)$					-0.012
					(0.017)
Rested		0.229^{***}		0.280^{***}	0.287^{***}
		(0.023)		(0.027)	(0.027)
MP			0.018^{***}	0.018^{***}	0.018^{***}
			(0.007)	(0.007)	(0.007)
$\hat{e}$			-0.043^{***}	-0.042^{***}	-0.042^{***}
			(0.007)	(0.007)	(0.007)
Constant	-1.899^{***}	-2.010^{***}	-2.088^{***}	-2.218^{***}	-2.227^{***}
	(0.045)	(0.045)	(0.077)	(0.081)	(0.079)
Marginal effects on Pr(injury), percentage points:					
WHP(1)	0.423^{***}	0.889^{***}	-0.958^{**}	-0.389	-0.326
Rested		2.676^{***}		2.724^{***}	2.843^{***}
MP			0.151^{***}	0.137^{***}	0.136^{***}
$1/\sigma$			1.052	1.050	1.050
IV F-Stat			594.749	593.331	581.631
Observations	160,615	160,615	160,615	160,615	160,615

Note: Marginal effects on the probability of injury, in percentage points, evaluated at a 10th-decile player in the base season with continuous regressors at their medians. For the binary regressor Rested, the effect is the discrete change in probability from 0 to 1. *p<0.1; **p<0.05; ***p<0.01.

Table 3: Probit estimates of the relationship between injury and one week chronic workload with decile and season fixed effects.

adding MP_{ig} to the model does not offset the negative bias arising from selection on persistence in injury risk, hence the negative coefficient when $Rested_{ig}$ is omitted and MP_{ig} is included and instrumented.

Next, we address all sources of omitted variable bias simultaneously, estimating the fully specified model (7) via probit-CF and obtaining a negative and statistically insignificant

coefficient for $WHP_{ig}(1)$. Taken at face value, this result indicates that one-week chronic workload has no effect on player injury risk once current assigned work and persistent injury risk shocks are controlled for, and it implies that the positive coefficient for $WHP_{ig}(1)$ in specification (1) is wholly spurious.

Finally, we check whether the relationship between injury and chronic workload could be nonlinear and U-shaped, as claimed in the literature. This would be the case if only extreme values of chronic workload generate a higher probability of injury. The probit model of interest is thus given by

$$\text{Injury}_{ig} = \mathbb{1}(\beta_0 + \beta_1 \text{WHP}_{ig}(X) + \beta_2 \text{WHP}_{ig}^2(X) + \beta_3 \text{Rested}_{ig} + \beta_4 \text{MP}_{ig} + \epsilon_{ig}) \tag{8}$$

where the chronic workload hypothesis implies $\beta_1 < 0$ and $\beta_2 > 0$. As seen in specification (5) of Table 3, the model yields negative and statistically insignificant coefficients for both $WHP_{ig}(1)$ and $WHP_{ig}^2(1)$. These coefficients are jointly insignificant and have signs that are inconsistent with the U-shaped relationship documented in the literature. All the other coefficients barely change. Our findings suggest *no relationship*—whether linear or U-shaped—between a player's chronic workload over a one-week horizon and their risk of injury, when controlling for current assigned work and persistent injury risk shocks.[17] Specifications using $WHP_{ig}(X)$ for $X = 2, 3, 4$, corresponding to two-, three-, and four-week horizons, produce very similar results (see Appendix B).

To sum up, our estimates indicate no partial relationship between player injury risk and chronic workload, whether linear or U-shaped, at any typical horizon. As our control for stress shocks is far from perfect, we should be careful not to interpret the coefficient on WHP_{ig} as causal. However, the specifications in Table 3 demonstrate that naive coefficient estimates for chronic workload are inherently unstable, and would become either more positive or more negative as different sources of endogeneity feature more or less prominently in the data-generating process and are controlled for, which likely explains the variety of results

[17] Intermediate specifications for equation (8) can be found in Appendix B.

obtained in the literature.

Regarding potential flaws in this analysis, one may be concerned that rest reduces players' conditioning and therefore raises the risk of injury [see for instance Gabbett (2016)]. In this case, the coefficient on $Rested_{ig}$ would partly capture the effect of reduced workload instead of serving as a control for prior stress shocks. This is a valid concern. In the next section, we present evidence to support the claim that $Rested_{ig}$, and the coach-selection component of $Rested_{ig}$ in particular, captures coach responses to stress shocks rather than the effect of rest on injury risk, and that rest itself does not cause an increase in injury risk.

5.3 Selection on Persistent Variation in Injury Risk

In this section, we examine whether rest itself causes injury and whether coaches effectively intervene on the extensive margin in response to player injury risk. The qualification "effectively" is important here because coaches demonstrably make such decisions, and strong evidence suggests that these decisions are often based on genuine beliefs about player injury risk.[18] What is at issue, then, is whether or not coaches accurately identify instances in which players are at heightened risk of injury. This is what we turn to next.

The key variable in the analysis is an indicator for instances in which a player received extra rest prior to the current game, which is similar to the variable $Rested_{ig}$ defined in Section 5.2. However, in this section, we define an extra-rest variable that distinguishes among the reasons for allocating the extra rest. In particular, let the variable $Rested_{ig}(X)$ be defined as

$$\text{Rested}_{ig}(X) \equiv \mathbb{1}(\text{MP}_{ig-1} = 0) \cdot \mathbb{1}(\text{Reason}_{ig-1} = X),$$

where $Reason_{ig-1}$ is the reason player i was rested during game $g-1$. In other words, the variable $Rested_{ig}(X)$ takes on a value of one if the player was rested during the previous game and if X is the reason the player was rested. The variable X has support

[18]Since the NBA implemented the Player Participation Policy, under which teams are fined for "illegal load management," coaches have violated the policy and incurred fines of $100,000 for doing so on several occasions, citing injury risk as their motivation and indicating a willingness to pay to offset the perceived risk. See for instance Vorkunov (2024).

$\{Recovery, Coach, Exogenous\}$, with its values defined as follows:

- *Recovery* denotes instances in which the player was out while recovering from a previous injury.
- *Coach* denotes instances in which players were available to participate (i.e., were not absent because of injury, illness, or personal/technical reasons) but the coach decided not to field them.
- *Exogenous* denotes instances in which players were out due to illness or personal/technical reasons.

If we set $X = \{Recovery, Coach\}$, then $Rested_{ig}(X)$ is precisely equal to the variable $Rested_{ig}$ in Section 5.2. Setting $X = Coach$ focuses exclusively on instances in which a coach decided not to field a player even though they were available and is the main case of interest in this section. Finally, setting $X = Exogenous$ allows us to estimate player injury risk after they missed a game because of circumstances unrelated to injury and beyond the control of coaches.

The key assumption we make in this section is that absence due to illness or personal/technical reasons is orthogonal to coach assessments of player injury risk. Hence, if extra rest causally increases injury risk because it affects player conditioning (as claimed by the literature), then both $X = Coach$ and $X = Exogenous$ should yield similar positive coefficients for $Rested_{ig}(X)$. However, if extra rest does not affect player conditioning and coach-assigned rest instead indicates that coaches are responding to persistent injury risk shocks, then $X = Coach$ should yield a positive coefficient for $Rested_{ig}(X)$ and $X = Exogenous$ should not.

To test these hypotheses, we estimate the following probit model for player injury:

$$\text{Injury}_{ig} = \mathbb{1}(\beta_0 + \beta_1 \text{MP}_{ig} + \beta_2 \text{Rested}_{ig}(X) + \beta_3 \text{MP}_{ig} \cdot \text{Rested}_{ig}(X) + \epsilon_{ig} > 0). \tag{9}$$

For each value of X, we estimate two versions of this model: one in which β_3, the coefficient

	Reason X for extra rest					
	Recovery or Coach		Coach		Exogenous	
	(1)	(2)	(3)	(4)	(5)	(6)
Rested(X)	0.297***	0.146***	0.071	−0.037	−0.108	−0.019
	(0.041)	(0.048)	(0.054)	(0.061)	(0.079)	(0.151)
MP	0.016***	0.014***	0.014***	0.013***	0.014***	0.014***
	(0.004)	(0.004)	(0.005)	(0.005)	(0.004)	(0.004)
MP·Rested(X)		0.010***		0.010***		−0.005
		(0.002)		(0.003)		(0.007)
$\hat{e}$	−0.033***	−0.032***	−0.033***	−0.033***	−0.031***	−0.031***
	(0.005)	(0.005)	(0.005)	(0.005)	(0.004)	(0.004)
Constant	−2.213***	−2.154***	−2.113***	−2.090***	−2.079***	−2.080***
	(0.085)	(0.084)	(0.089)	(0.089)	(0.073)	(0.074)
Marginal effects on Pr(injury), percentage points:						
Rested(X)	2.987***	4.117***	0.613	2.060***	−0.808	−0.928*
MP	0.124***		0.117***		0.114***	
MP | Rested(X)=0		0.109***		0.110***		0.114***
MP | Rested(X)=1		0.342***		0.277***		0.061
$1/\sigma$	1.039	1.037	1.040	1.038	1.037	1.037
IV F-Stat	605.518	605.518	583.497	583.497	607.196	607.196
Observations	160,615	160,615	160,615	160,615	160,615	160,615

Note: Marginal effects on the probability of injury, in percentage points, evaluated at a 10th-decile player in the base season with continuous regressors at their medians. For the binary regressor Rested(X), the effect is the discrete change in probability from 0 to 1. *p<0.1; **p<0.05; ***p<0.01.

Table 4: Test for the direction of causality in the relationship between extra rest and player injury risk.

on the interaction term $MP_{ig} \cdot Rested_{ig}(X)$, is restricted to zero, and another in which β_3 is allowed to vary freely. In addition to level differences in player injury risk, the latter specification is set up to capture differences in the marginal risk of injury after rest with respect to current work, MP_{ig}. Including the interaction term $MP_{ig} \cdot Rested_{ig}(X)$ allows the model to account for differences in current work conditional on $Rested_{ig}(X)$ and to adjust the predicted risk of injury accordingly.

Table 4 shows estimates of the parameters of model (9) with and without the interaction term $MP_{ig} \cdot Rested_{ig}(X)$. Every specification includes salary-decile and season fixed effects, with heteroskedasticity-robust standard errors clustered at the player-season level and adjusted for the generated regressor $\hat{e}_{ig}$. As in the preceding tables, the panel beneath the coefficient estimates reports the implied marginal effects on the probability of injury, in percentage points; for the interacted specifications, the marginal effect of MP_{ig} is shown separately at $Rested_{ig}(X) = 0$ and $Rested_{ig}(X) = 1$.

Specification (1) shows the coefficient for $Rested_{ig}(X)$ when $X = \{Recovery, Coach\}$—which again is numerically identical to the variable $Rested_{ig}$ defined in Section 5.2—when the interaction with MP_{ig} is excluded. Just as in Section 5.2, $Rested_{ig}(X)$ has a large, positive, and statistically significant coefficient. Specification (2) expands on (1) by adding the interaction between $Rested_{ig}(X)$ and MP_{ig}. Here we again see a positive coefficient for $Rested_{ig}(X)$ as well as a positive and statistically significant coefficient for its interaction with MP_{ig}.

In specification (3), we set $X = Coach$ so that $Rested_{ig}(X)$ captures only instances in which a player was available in the previous game but did not play because the coach decided not to field them. The coefficient for $Rested_{ig}(X)$ in this specification is positive but is not statistically significant at any conventional level. However, when the interaction term $MP \cdot Rested_{ig}(X)$ is added in specification (4), its coefficient is positive, statistically significant, and very similar to its $X = \{Recovery, Coach\}$ counterpart. In other words, players seem to exhibit a lower baseline risk of injury after rest when rest due to injuries is excluded from $Rested_{ig}(X)$ but exhibit a similar marginal risk of injury.

Simply comparing the coefficient magnitudes for $Rested_{ig}(X)$ and its interaction gives some indication that the marginal risk of injury varies substantially with $Rested_{ig}(X)$, but it is worth briefly exploring the estimated marginal risk of injury implied by specification (4). Specifically, if we again consider 10th-decile players in the base season, specification (4) implies a marginal effect of 0.0011 with respect to MP_{ig} when it is at its median and $Rested_{ig}(X) = 0$. By contrast, for $Rested_{ig}(X) = 1$, the marginal effect with respect to MP_{ig}

is 0.0028 for the same subpopulation of players, which represents a 152% increase in the marginal probability of injury.

Finally, specification (5) shows the coefficient for $Rested_{ig}(X)$ obtained for $X = Exogenous$, excluding the interaction with MP_{ig}. The coefficient on $Rested_{ig}(X)$ is statistically insignificant and, in this instance, is also negative. Moreover, adding the interaction term again yields negative but statistically insignificant coefficients for both $Rested_{ig}(X)$ and its interaction with MP_{ig}. Thus, neither specification implies an increase in injury risk due to rest when that rest is not due to injury or endogenous decisions by coaches.

Overall, the estimates presented in Table 4 support the hypothesis that coaches respond to persistent injury risk shocks by resting players, and that rest itself does not increase a player's injury risk.

6 Conclusion

We began by asking whether managers hold and act on private, time-varying information about which worker is at risk. Three answers follow. First, once we instrument realized minutes with teammate absence, the effect of current work on injury is positive and significant. The negative association reported in the sports-medicine literature is an artifact of injury cutting play short and coaches sparing the vulnerable, not a sign that work protects players. Second, holding current assigned work and recent absence fixed, no relationship survives between injury and chronic workload, linear or U-shaped, at horizons of one to four weeks. Third, and most important, injury risk rises after a coach rests an available player but not after an illness or personal absence. Coaches are responding to something we do not see, not causing injury by over-resting.

These results speak first to the economics of management. Managers here act on private, hard-to-verify information about a worker's *state*, a signal no public personnel record captures. Ignoring how they assign work is enough to invert the measured relationship between work and injury. The mechanism is not special to basketball. Wherever the best workers are

skilled and hard to replace, a manager who can pull back on the vulnerable can raise output tomorrow by lowering it today. That logic cautions against blunt limits on managerial discretion, which forbid exactly the informed pullback we find coaches exercising.

The clearest instance of such a limit is already in place. Beginning in the 2023-24 season, just after our sample ends, the NBA's Player Participation Policy fines teams for resting healthy stars. In the terms of our analysis, it severs one link in the mechanism we document: the coach's ability to act on his private assessment of who is at risk. If that assessment is as informative as our estimates imply, constraining it should leave more vulnerable players on the court and produce more injuries, not fewer. Whether it does, and at what cost to the very availability the policy was meant to protect, is the object of further study.

References

Abaluck, J., Agha, L., Kabrhel, C., Raja, A. and Venkatesh, A. (2016). The determinants of productivity in medical testing: Intensity and allocation of care, *American Economic Review* **106**(12): 3730–3764.

Adhvaryu, A., Kala, N. and Nyshadham, A. (2022). Management and shocks to worker productivity, *Journal of Political Economy* **130**(1): 1–47.

Adhvaryu, A., Nyshadham, A. and Tamayo, J. (2023). Managerial quality and productivity dynamics, *The Review of Economic Studies* **90**(4): 1569–1607.

Amodio, F. and Martinez-Carrasco, M. A. (2018). Input allocation, workforce management and productivity spillovers: Evidence from personnel data, *The Review of Economic Studies* **85**(4): 1937–1970.

Artz, B. and Heywood, J. S. (2015). Performance pay and workplace injury: Panel evidence, *Economica* **82**(S1): 1241–1260.

Bache-Mathiesen, L. K., Andersen, T. E., Dalen-Lorentsen, T., Clarsen, B. and Fagerland, M. W. (2021). Not straightforward: Modelling non-linearity in training load and injury research, *BMJ Open Sport & Exercise Medicine* **7**(3): e001119.

Bache-Mathiesen, L. K., Andersen, T. E., Dalen-Lorentsen, T., Clarsen, B. and Fagerland, M. W. (2022). Assessing the cumulative effect of long-term training load on the risk of injury in team sports, *BMJ Open Sport & Exercise Medicine* **8**(2): e001342.

Baktash, M. B., Heywood, J. S. and Jirjahn, U. (2022). Worker stress and performance pay: German survey evidence, *Journal of Economic Behavior & Organization* **201**: 276–291.

Berniell, I. and Bietenbeck, J. (2020). The effect of working hours on health, *Economics & Human Biology* **39**: 100901.

Bloom, N., Eifert, B., Mahajan, A., McKenzie, D. and Roberts, J. (2013). Does management matter? Evidence from India, *The Quarterly Journal of Economics* **128**(1): 1–51.

Bullock, G. S., Ferguson, T., Arundale, A. H., Martin, C. L., Collins, G. S. and Kluzek, S. (2022). Return to performance following severe ankle, knee, and hip injuries in National Basketball Association players, *PNAS Nexus* **1**(4): 1–7.

Burgess, S., Propper, C., Ratto, M., Kessler Scholder, S. v. H. and Tominey, E. (2010). Smarter task assignment or greater effort: The impact of incentives on team performance, *The Economic Journal* **120**(547): 968–989.

Caparrós, T., Alentorn-Geli, E., Myer, G. D., Capdevila, L., Samuelsson, K., Hamilton, B. and Rodas, G. (2016). The relationship of practice exposure and injury rate on game performance and season success in professional male basketball, *Journal of Sports Science & Medicine* **15**(3): 397–402.

Caparrós, T., Casals, M., Solana, Á. and Peña, J. (2018). Low external workloads are related to higher injury risk in professional male basketball games, *Journal of Sports Science & Medicine* **17**(2): 289–297.

Cross, M. J., Williams, S., Trewartha, G., Kemp, S. P. and Stokes, K. A. (2016). The influence of in-season training loads on injury risk in professional rugby union, *International Journal of Sports Physiology and Performance* **11**(3): 350–355.

Currie, J. and MacLeod, W. B. (2017). Diagnosing expertise: Human capital, decision making, and performance among physicians, *Journal of Labor Economics* **35**(1): 1–43.

Cygan-Rehm, K. and Wunder, C. (2018). Do working hours affect health? Evidence from statutory workweek regulations in Germany, *Labour Economics* **53**: 162–171.

Fulton, J., Wright, K., Kelly, M., Zebrosky, B., Zanis, M., Drvol, C. and Butler, R. (2014). Injury risk is altered by previous injury: A systematic review of the literature and presentation of causative neuromuscular factors, *International Journal of Sports Physical Therapy* **9**(5): 583–595.

Gabbett, T. J. (2016). The training-injury prevention paradox: Should athletes be training smarter and harder?, *British Journal of Sports Medicine* **50**(5): 273–280.

Gabbett, T. J. and Ullah, S. (2012). Relationship between running loads and soft-tissue injury in elite team sport athletes, *The Journal of Strength & Conditioning Research* **26**(4): 953–960.

Griffin, A., Kenny, I. C., Comyns, T. M. and Lyons, M. (2020). The association between the acute:chronic workload ratio and injury and its application in team sports: A systematic review, *Sports Medicine* **50**(3): 561–580.

Hansen, B. E. and Lee, S. (2019). Asymptotic theory for clustered samples, *Journal of Econometrics* **210**(2): 268–290.

Hoffman, M., Kahn, L. B. and Li, D. (2018). Discretion in hiring, *The Quarterly Journal of Economics* **133**(2): 765–800.

Hoffman, M. and Tadelis, S. (2021). People management skills, employee attrition, and manager rewards: An empirical analysis, *Journal of Political Economy* **129**(1): 243–285.

Hulin, B. T., Gabbett, T. J., Blanch, P., Chapman, P., Bailey, D. and Orchard, J. W. (2014). Spikes in acute workload are associated with increased injury risk in elite cricket fast bowlers, *British Journal of Sports Medicine* **48**(8): 708–712.

Impellizzeri, F. M., Tenan, M. S., Kempton, T., Novak, A. and Coutts, A. J. (2020). Acute:chronic workload ratio: Conceptual issues and fundamental pitfalls, *International Journal of Sports Physiology and Performance* **15**(6): 907–913.

Impellizzeri, F. M., Woodcock, S., Coutts, A. J., Fanchini, M., McCall, A. and Vigotsky, A. D. (2021). What role do chronic workloads play in the acute to chronic workload ratio? Time to dismiss ACWR and its underlying theory, *Sports Medicine* **51**(3): 581–592.

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

Kraemer, W. J., Ratamess, N. A. and French, D. N. (2002). Resistance training for health and performance, *Current Sports Medicine Reports* **1**(3): 165–171.

Lazear, E. P. (2000). Performance pay and productivity, *American Economic Review* **90**(5): 1346–1361.

Lazear, E. P., Shaw, K. L. and Stanton, C. T. (2015). The value of bosses, *Journal of Labor Economics* **33**(4): 823–861.

Lee, J. and Lee, Y.-K. (2016). Can working hour reduction save workers?, *Labour Economics* **40**: 25–36.

Lewis, M. (2018). It's a hard-knock life: Game load, fatigue, and injury risk in the National Basketball Association, *Journal of Athletic Training* **53**(5): 503–509.

Marousek, F. (2005). Pro Sports Transactions Archive, ProSportsTransactions.com. Accessed: April 14, 2025.
URL: *https://www.prosportstransactions.com/*

Menon, S., Morikawa, L., Tummala, S. V., Buckner-Petty, S. and Chhabra, A. (2024). The primary risk factors for season-ending injuries in professional basketball are minutes played per game and later season games, *Arthroscopy: The Journal of Arthroscopic & Related Surgery* **40**(9): 2468–2473.

Milanez, A., Lemmens, A. and Ruggiu, C. (2025). Algorithmic management in the workplace: New evidence from an OECD employer survey, *OECD Artificial Intelligence Papers 31*, OECD Publishing, Paris.

Murphy, K. M. and Topel, R. H. (2002). Estimation and inference in two-step econometric models, *Journal of Business & Economic Statistics* **20**(1): 88–97.

National Basketball Association and National Basketball Players Association (2023). Collective Bargaining Agreement. Article XXII (Player Health and Wellness). Accessed: August 6, 2026. Available at https://ak-static.cms.nba.com/wp-content/uploads/sites/4/2023/06/2023-NBA-Collective-Bargaining-Agreement.pdf.

Nekoei, A., Sigurdsson, J. and Wehr, D. (2025). The economic burden of burnout, *CESifo Working Paper 11128*, CESifo, Munich.

Orchard, J. W., James, T., Portus, M., Kountouris, A. and Dennis, R. (2009). Fast bowlers in cricket demonstrate up to 3-to 4-week delay between high workloads and increased risk of injury, *The American Journal of Sports Medicine* **37**(6): 1186–1192.

Powell, S. (2024). NBA sends teams the results found in load management study. *NBA.com*. Accessed: August 9, 2026. Available at https://www.nba.com/news/nba-sends-data-load-management-study.

Prasad, K., McLoughlin, C., Stillman, M., Poplau, S., Goelz, E., Taylor, S., Nankivil, N., Brown, R., Linzer, M., Cappelucci, K., Barbouche, M. and Sinsky, C. A. (2021). Prevalence and correlates of stress and burnout among U.S. healthcare workers during the COVID-19 pandemic: A national cross-sectional survey study, *eClinicalMedicine* **35**: 100879.

Prendergast, C. (1999). The provision of incentives in firms, *Journal of Economic Literature* **37**(1): 7–63.

Rivers, D. and Vuong, Q. H. (1988). Limited information estimators and exogeneity tests for simultaneous probit models, *Journal of Econometrics* **39**(3): 347–366.

Soligard, T., Schwellnus, M., Alonso, J.-M., Bahr, R., Clarsen, B., Dijkstra, H. P., Gabbett, T., Gleeson, M., Hägglund, M., Hutchinson, M. R., Janse van Rensburg, C., Khan, K. M., Meeusen, R., Orchard, J. W., Pluim, B. M., Raftery, M., Budgett, R. and Engebretsen, L. (2016). How much is too much? (Part 1) International Olympic Committee consensus statement on load in sport and risk of injury, *British Journal of Sports Medicine* **50**(17): 1030–1041.

Toohey, L. A., Drew, M. K., Cook, J. L., Finch, C. F. and Gaida, J. E. (2017). Is subsequent lower limb injury associated with previous injury? A systematic review and meta-analysis, *British Journal of Sports Medicine* **51**(23): 1670–1678.

Trudell, M. (2018). A look inside the Lakers training staff, Los Angeles Lakers, NBA.com. Accessed: August 2, 2026.
URL: *https://www.nba.com/lakers/news/180225-lakers-training-staff*

Vorkunov, M. (2024). Nets first to be fined under NBA's new player rest policy after sitting 4 versus Bucks, The Athletic. Accessed: April 18, 2025.
URL: *https://www.nytimes.com/athletic/5179459/2024/01/04/nets-fine-nba-player-rest-policy/*

Wong, J., Whatman, C., Sheerin, K., Borotkanics, R. and Wells, D. (2022). The relationship between training and game workload and injury risk in elite male basketball, *New Zealand Journal of Sports Medicine* **49**(1): 10–15.

Wooldridge, J. M. (2010). *Econometric Analysis of Cross Section and Panel Data*, 2 edn, MIT press, Cambridge, Massachusetts.

Wooldridge, J. M. (2015). Control function methods in applied econometrics, *Journal of Human Resources* **50**(2): 420–445.

A Notes on Data Cleaning

To ensure accuracy with respect to the onset and duration of injuries, we cross-checked the ProSportsTransactions.com (PST) absenteeism data—which record absences due to injury, illness, and personal and technical reasons—with each player's status as reported in the NBA.com official box-score data (NDC). On game days, the latter report the inactive list and, for players who were active but did not play, stipulate whether this was due to injury, illness, suspension, personal reasons, or the coach's decision. From this dataset of game-day instances, we constructed another event-level dataset documenting time intervals (i.e., start and end dates) during which players were on the inactive list and/or missed one or more games. We then matched this dataset to the PST event-level dataset twice, once by player and event start date and again by player and event end date. Of the $7,816$ injury-list and missed-game events recorded in PST, $6,619$ (84.7%) matched the NDC data on end dates and $6,012$ (76.9%) matched on start dates. Collectively, $7,458$ (95.4%) matched on a start date, an end date, or both, and $4,847$ (62.0%) matched a consistent event on both start and end dates.[19]

For 358 events in PST that did not match any periods of inactivity in the NDC data, we looked for NDC events whose date intervals overlapped the PST events and corresponded to inactive-list status, and found a unique match for 331 observations. PST and NDC dates typically differ by a day at each boundary: PST records the date of formal inactive-list movement, while NDC reports the first and last missed games. Among the remaining 27 unmatched observations, another 12 were found to correspond to a contiguous NDC inactive interval that enclosed the PST start date, the PST end date, or both—often carrying a different inactivity label (e.g., coach's decision or personal) from the PST injury label. The remaining 15 observations involved season-long absences, events that occurred entirely between games or outside the season, or events that coincided with a G League assignment.

[19]Of those matched on both boundary dates, 326 matched distinct NDC events on the start- and end-date sides rather than a single event; we treat these as partial matches and address them in Cases 1–4 below.

The rest of the sample consists of partial matches on end dates and start dates, with the unmatched end or start date in PST either falling before or after the corresponding date for the proposed match in NDC. Thus, we broke the partially matched data into four cases based first on whether the partial match was between start dates or end dates and then on whether the unmatched date in PST fell before or after the corresponding date for the proposed match in NDC. We treat each of these cases in turn below.

Case 1: First, we examined the 381 partial matches on start dates in which the PST end date fell after the NDC end date. Among these, we found 165 instances in which the player was active on dates falling between the two event end dates. Inspection of the events revealed that these were observations in which the PST start and end dates were obtained from two distinct events. In each of these instances, an event end date was missing in PST, leading to a merging of events. We split these events using start and end dates from NDC to complete the PST event intervals. Generated events corresponding to a period in which the player was active were dropped. Splitting these 165 events yielded a total of 265 inactive events, which were merged back into the PST event dataset after the events from which they were originally generated were dropped. Upon completion of this process, 217 events remained in which the PST end date was later than the matched NDC end date—the 216 originals not affected by the split, plus one split-derived event whose newly matched NDC end date still preceded the recorded PST end date. Of these observations, 54 corresponded to games missed because of a G League relegation or for personal or disciplinary reasons. Of the remaining 163 events, only one had an active game between the two end dates, and the NDC date was used for this event. For the remaining 162 events, the PST end date was assumed to be correct, since there were no active games between the PST and NDC end dates.

Case 2: Next, we examined the $1,127$ partial start-date matches in which the PST end date occurred before the NDC end date. With one exception, players were inactive in all games between these dates. However, it was important to determine whether the PST end date was the actual date on which a player recovered from their injury so as to correctly

capture injury durations from the data. To this end, we examined the data to find reasons for end-date discrepancies. First, in 55 instances, we found that the player was traded, waived, or retired a few days after their PST end date. For these observations, activation in PST occurred mechanically as part of this process and thus censored the true recovery dates. Of the remaining 1,072 observations, 804 had no games between the PST and NDC end dates, implying no contradiction between the two datasets and a reliable PST end date. In a further 22, the player had no subsequent active game in our sample and none of the patterns enumerated below applied; with no NBA.com games against which to adjudicate the end date, we retained the PST end date for these observations. This left 246 questionable observations. Of these, 26 were followed by an entire additional injury event, already recorded elsewhere in the PST event dataset, that accounted for the additional inactive games. Another 70 were found to have two activation dates recorded in PST for the same injury, one indicating a movement off the inactive list and the other indicating a return to play. In these cases, we took the latter date to be correct, which corresponded to the date in NDC. Among the remaining observations, 67 were instances in which the player spent additional time out for disciplinary or personal reasons or was relegated to the G League before returning to the active roster. This left 83 observations for which we could not find a systematic reason for the discrepancy in end dates. We examined each of these individually, comparing them with injury reports at FoxSports.com regarding the player's recovery, and found that the NDC end date was correct in 38 instances. Of these, 29 were instances in which the two end dates differed by just two days. For the rest, the PST end date was found to be correct.

Case 3: We examined the 290 partial end-date matches in which the NDC start date preceded the PST start date, which posed fewer issues than the partial start-date matches in Cases 1–2. In 83 instances, we found an alternative, noninjury explanation for the early NDC start date: a G League assignment, personal leave, or disciplinary suspension recorded at some point between the two start dates. We therefore retained the PST start date. Of the remaining 207 observations, 105 had no games between the two start dates, implying no contradiction with the PST start date. This left 102 observations in which there was at least

one game between the two start dates. For 21 of these, the games were attributable to a preceding event already captured in the PST event-level dataset, so we again retained the PST start date. For the remaining 81, we cross-checked against injury reports at FoxSports.com and found 12 instances in which the PST start date needed to be adjusted, in most cases by only a few days. The remaining 69 observations were retained.

Case 4: Finally, after correcting for 34 instances in which an injury was recorded on the day it occurred instead of the day after, there were $1,677$ instances in which the PST start date preceded the NDC start date. However, $1,455$ of these had no games between the two dates. Among the remaining 222 observations, 62 were instances in which a player was active between the two start dates. After verifying a sample of these by hand, we concluded that the matched end date was coincidental and reset it to the day before the next game in which the player was active according to NBA.com per-game records; this is typically a non-game calendar day rather than the last observed inactive day. Two further observations had a PST start preceding the regular-season opening with no intervening active game and required no correction. This left 158 of the 222 questionable observations. For those with the smallest start-date gaps (up to roughly four days), we cross-checked against FoxSports.com injury reports and found no cases that needed correction; cases with larger gaps were retained as recorded.

Upon completion of this process, we had a total of $7,916$ events. However, upon checking the data, we found 65 player-specific events with shared end dates. We examined these on a case-by-case basis to eliminate conflicts. Of these, 2 observations were exact duplicates of another record for the same player, with the same start date, end date, and description, and we dropped the redundant entry. A further 41 observations were part of groups of records describing the same underlying injury but with different start dates or different status descriptions. In each group, we kept the event whose start date we verified against external sources, typically the earliest event in the group, and dropped the rest. The remaining 22 observations were instances in which one injury event was followed by another before the player returned to play. For example, a player who was out due to health-and-safety

protocols also sometimes spent time conditioning before returning to play, and these periods were recorded as separate events with the same end date. In these cases, we set the end date of the earlier event to fall before the start of the next event, typically on the preceding day. In total, we dropped 25 observations to obtain a trimmed sample of $7,891$ intervals corresponding to inactive-list movements and games missed due to injury. We then checked that all remaining events were chronologically disjoint, dropping two events contained entirely within another and correcting a surgical-recovery date by one day to obtain $7,889$ disjoint event observations.

One of the key weaknesses of publicly available injury data is incompleteness. To fill some of the gaps in the PST injury data, we added injury and illness information from a third source, FoxSports.com. In the Fox Sports (FS) data, injuries are recorded with a brief description and a date—i.e., FS provides no duration information. To incorporate these data, we first looked for matches between FS injuries and those in PST by matching the FS date with the PST start date via the following iterative process. First, we looked for exact matches between the FS date and the PST start date, recorded them, and then removed the matched observations from FS. Thereafter, we repeated this process between the FS date and dates one, two, three, and four days before the PST start date, iteratively removing the matched FS observations to create a sample of unmatched observations after each attempted match. We then took the remaining unmatched observations and repeated this process with dates one, two, three, and four days after the PST start date. This process yielded

- $3,352$ same-day matches;
- $1,805$ matches one day before, 332 matches two days before, 74 matches three days before, and 30 matches four days before; and
- 230 matches one day after, 87 matches two days after, 37 matches three days after, and 25 matches four days after.

We then repeated this process on the full FS sample using the NDC event data and obtained:

- $2,982$ same-day matches;

- 2,557 matches one day before, 627 matches two days before, 186 matches three days before, and 85 matches four days before; and
- 201 matches one day after, 100 matches two days after, 47 matches three days after, and 49 matches four days after.

When an FS event matches a PST event, it indicates that the event has already been recorded in the injury dataset. Excluding these events, we retained 1,028 FS-NDC matches, comprising:

- 439 same-day matches;
- 344 matches one day before, 43 matches two days before, 45 matches three days before, and 25 matches four days before; and
- 61 matches one day after, 29 matches two days after, 17 matches three days after, and 25 matches four days after.

For this sample of matched FS events, we recorded the last date of the NDC event as the FS event end date. We then looked for instances in which FS-NDC matches overlapped completely or partially with events in PST to ensure that none of the FS-NDC matches duplicated previously recorded events. This flagged 84 duplicate events. Removing these yielded 944 additional unique events obtained from the FS-NDC matching process.

Finally, we constructed a consolidated event dataset initially comprising 8,833 injury and illness events. We then matched this consolidated event dataset with the NDC box-score data to ensure that no event had a start date corresponding to a day on which the player was active. (Event start dates were coded to correspond to the day on which the player's absence started.) We found just 6 violations, which were examined and fixed on a case-by-case basis. A final chronology pass dropped events that overlapped an earlier event for the same player, yielding 8,812 injury and illness events.

In matching the FS events, we assumed that FS tended to report injuries as they occurred rather than as they affected team lineups. This is because the FS injury reports summarize

news articles about injury events, whereas PST records injury-list movements. We therefore matched FS events with PST and NDC events by looking first for matches on days preceding the PST and NDC start dates and considering days after those start dates only once the earlier matches had been removed. The higher share of matches found between FS dates and dates just preceding the PST and NDC start dates partly corroborated our assumption. As a robustness check, we reran the matching with the offset order reversed. The concentration of matches just before the start date remained, though it was somewhat weaker.

With the combined PST-FS event dataset in hand, we proceeded to match these events with games in the NDC box-score data. Recall that the event-level dataset records events with a start date corresponding to the beginning of the inactivity event in the NDC box-score data. Thus, the game in the NDC box-score data that matches the PST-FS event is the game immediately preceding the event. In some instances, this implied a match between an injury event and a game in which a player did not participate. While this could indicate an erroneously late start date in the event data, most instances occurred among rookies and low-paid players, who are absent for a greater share of games in general and often spend significant time playing in the NBA G League. Thus, these discrepancies could plausibly be attributed to injuries outside official NBA games. Finally, PST and FS notes were cross-checked to classify player absenteeism as being due to a play-related injury, an illness, or personal/technical reasons.

B Additional Reduced-Form Specifications

B.1 Probit-CF chronic workload estimates at two-, three-, and four-week horizons

This section reports probit-CF estimates. All specifications include salary-decile and season fixed effects. All standard errors are clustered at the player-season level, are robust to heteroskedasticity, and have been adjusted to account for the generated regressor $\hat{e}$. (See Appendix C for the derivation of these standard errors.)

	(1)	(2)	(3)	(4)
WHP(2)	0.046^{***}	0.089^{***}	-0.115^{**}	-0.063
	(0.014)	(0.014)	(0.057)	(0.054)
Rested		0.212^{***}		0.275^{***}
		(0.022)		(0.029)
MP			0.017^{***}	0.017^{***}
			(0.006)	(0.006)
$\hat{e}$			-0.042^{***}	-0.042^{***}
			(0.006)	(0.006)
Constant	-1.899^{***}	-2.000^{***}	-2.073^{***}	-2.205^{***}
	(0.045)	(0.045)	(0.073)	(0.078)
$1/\sigma$			1.048	1.048
IV F-Stat			719.104	708.552
Observations	160,615	160,615	160,615	160,615
Note:				*p<0.1; **p<0.05; ***p<0.01

Table 5: Probit estimates of the linear relationship between injury and two week chronic workload with decile and season fixed effects.

	(1)	(2)	(3)	(4)
WHP(2)	0.005	0.148^{***}	-0.204^{**}	-0.021
	(0.042)	(0.043)	(0.083)	(0.075)
$\text{WHP}^2(2)$	0.020	-0.027	0.044^{**}	-0.017
	(0.019)	(0.019)	(0.021)	(0.020)
Rested		0.221^{***}		0.281^{***}
		(0.022)		(0.028)
MP			0.017^{***}	0.017^{***}
			(0.006)	(0.006)
$\hat{e}$			-0.042^{***}	-0.042^{***}
			(0.006)	(0.006)
Constant	-1.890^{***}	-2.018^{***}	-2.052^{***}	-2.216^{***}
	(0.047)	(0.047)	(0.071)	(0.078)
$1/\sigma$			1.049	1.048
IV F-Stat			702.565	698.241
Observations	160,615	160,615	160,615	160,615
Note:				*p<0.1; **p<0.05; ***p<0.01

Table 6: Probit estimates of the non-linear relationship between injury and two week chronic workload with decile and season fixed effects.

	(1)	(2)	(3)	(4)
WHP(3)	0.045^{***}	0.083^{***}	-0.105^{**}	-0.064
	(0.014)	(0.014)	(0.053)	(0.050)
Rested		0.201^{***}		0.275^{***}
		(0.022)		(0.030)
MP			0.016^{***}	0.017^{***}
			(0.006)	(0.006)
$\hat{e}$			-0.040^{***}	-0.041^{***}
			(0.006)	(0.006)
Constant	-1.898^{***}	-1.991^{***}	-2.064^{***}	-2.196^{***}
	(0.045)	(0.045)	(0.070)	(0.077)
$1/\sigma$			1.045	1.046
IV F-Stat			808.385	790.574
Observations	160,615	160,615	160,615	160,615
Note:				*p<0.1; **p<0.05; ***p<0.01

Table 7: Probit estimates of the linear relationship between injury and three week chronic workload with decile and season fixed effects.

	(1)	(2)	(3)	(4)
WHP(3)	0.017	0.137***	−0.166**	−0.010
	(0.046)	(0.045)	(0.076)	(0.068)
$WHP^2(3)$	0.014	−0.026	0.031	−0.025
	(0.021)	(0.021)	(0.022)	(0.021)
Rested		0.207***		0.281***
		(0.022)		(0.030)
MP			0.016***	0.017***
			(0.006)	(0.006)
$\hat{e}$			−0.040***	−0.041***
			(0.006)	(0.006)
Constant	−1.892***	−2.006***	−2.051***	−2.209***
	(0.047)	(0.047)	(0.070)	(0.077)
$1/\sigma$			1.046	1.046
IV F-Stat			795.583	784.909
Observations	160,615	160,615	160,615	160,615
Note:				*p<0.1; **p<0.05; ***p<0.01

Table 8: Probit estimates of the non-linear relationship between injury and three week chronic workload with decile and season fixed effects.

	(1)	(2)	(3)	(4)
WHP(4)	0.050^{***}	0.084^{***}	-0.088^{*}	-0.054
	(0.014)	(0.014)	(0.050)	(0.047)
Rested		0.197^{***}		0.278^{***}
		(0.022)		(0.031)
MP			0.015^{***}	0.016^{***}
			(0.005)	(0.005)
$\hat{e}$			-0.039^{***}	-0.040^{***}
			(0.005)	(0.005)
Constant	-1.899^{***}	-1.988^{***}	-2.061^{***}	-2.195^{***}
	(0.045)	(0.045)	(0.069)	(0.076)
$1/\sigma$			1.044	1.045
IV F-Stat			850.83	830.072
Observations	160,615	160,615	160,615	160,615
Note:				*p<0.1; **p<0.05; ***p<0.01

Table 9: Probit estimates of the linear relationship between injury and four week chronic workload with decile and season fixed effects.

	(1)	(2)	(3)	(4)
WHP(4)	0.028	0.135^{***}	-0.133^{*}	0.004
	(0.047)	(0.046)	(0.071)	(0.064)
$\text{WHP}^2(4)$	0.011	-0.025	0.024	-0.028
	(0.022)	(0.022)	(0.023)	(0.022)
Rested		0.202^{***}		0.284^{***}
		(0.021)		(0.031)
MP			0.015^{***}	0.016^{***}
			(0.005)	(0.005)
$\hat{e}$			-0.039^{***}	-0.040^{***}
			(0.005)	(0.005)
Constant	-1.895^{***}	-2.001^{***}	-2.052^{***}	-2.209^{***}
	(0.047)	(0.047)	(0.069)	(0.077)
$1/\sigma$			1.044	1.045
IV F-Stat			840.641	827.592
Observations	160,615	160,615	160,615	160,615
Note:				*p<0.1; **p<0.05; ***p<0.01

Table 10: Probit estimates of the non-linear relationship between injury and four week chronic workload with decile and season fixed effects.

B.2 LPM-CF injury-workload estimates with salary-decile and season fixed effects

The models in this section are linear probability model (LPM) counterparts to the probit and probit-CF estimates presented in the main text. Except for the specifications in Table 11, all specifications include salary-decile and season fixed effects. All standard errors are clustered at the player-season level and are robust to heteroskedasticity. The stacked generated-regressor adjustment described in Appendix C applies only to the LPM-CF estimates in columns (5)–(6) of Table 11 and columns (2), (4), and (6) of Table 20; these covariance estimates also use the HC1-style multiplier $C/(C-1) \times N/(N-K)$. The remaining columns report clustered OLS or 2SLS standard errors and contain no generated regressor. For readability, all coefficients and standard errors are scaled by 100.

	OLS	OLS	2SLS	2SLS	LPM-CF	LPM-CF
	(1)	(2)	(3)	(4)	(5)	(6)
MP	-0.046^{***}	-0.116^{***}	0.110^{*}	0.107^{***}	0.047	0.084^{**}
	(0.006)	(0.007)	(0.066)	(0.031)	(0.068)	(0.036)
MP^2					0.002^{***}	0.001
					(0.0004)	(0.0004)
$\hat{e}$					-0.167^{**}	-0.233^{***}
					(0.067)	(0.032)
Constant	4.329^{***}		0.761		1.120	
	(0.153)		(1.510)		(1.514)	
Decile & Season FE		✓		✓		✓
IV F-Stat			69.24	603.38	69.24	603.38
Observations	160,615	160,615	160,615	160,615	160,615	160,615
Note:						*p<0.1; **p<0.05; ***p<0.01

Table 11: LPM estimates of risk of injury with respect to time on-court with decile and season fixed effects.

	(1)	(2)	(3)	(4)
WHP(1)	0.304***	0.663***	−0.881**	−0.430
	(0.089)	(0.093)	(0.399)	(0.375)
Rested		1.704***		2.005***
		(0.188)		(0.215)
MP			0.145***	0.142***
			(0.048)	(0.047)
IV F-Stat			594.749	593.331
Observations	160,615	160,615	160,615	160,615
Note:				*p<0.1; **p<0.05; ***p<0.01

Table 12: LPM estimates of the linear relationship between injury and one week chronic workload with decile and season fixed effects.

	(1)	(2)	(3)	(4)
WHP(1)	-0.225	0.951^{***}	-1.838^{***}	-0.452
	(0.255)	(0.268)	(0.573)	(0.526)
$WHP^2(1)$	0.252^{**}	-0.130	0.448^{***}	0.009
	(0.118)	(0.121)	(0.132)	(0.128)
Rested		1.770^{***}		2.000^{***}
		(0.195)		(0.212)
MP			0.147^{***}	0.142^{***}
			(0.048)	(0.047)
IV F-Stat			578.967	581.631
Observations	160,615	160,615	160,615	160,615
Note:				*p<0.1; **p<0.05; ***p<0.01

Table 13: LPM estimates of the non-linear relationship between injury and one week chronic workload with decile and season fixed effects.

	(1)	(2)	(3)	(4)
WHP(2)	0.336^{***}	0.650^{***}	-0.909^{**}	-0.554
	(0.099)	(0.101)	(0.410)	(0.386)
Rested		1.594^{***}		1.996^{***}
		(0.186)		(0.226)
MP			0.137^{***}	0.141^{***}
			(0.044)	(0.044)
IV F-Stat			719.104	708.552
Observations	160,615	160,615	160,615	160,615
Note:				*p<0.1; **p<0.05; ***p<0.01

Table 14: LPM estimates of the linear relationship between injury and two week chronic workload with decile and season fixed effects.

	(1)	(2)	(3)	(4)
WHP(2)	−0.007	1.037***	−1.675***	−0.443
	(0.302)	(0.308)	(0.588)	(0.527)
WHP^2(2)	0.170	−0.186	0.375**	−0.053
	(0.145)	(0.146)	(0.155)	(0.148)
Rested		1.657***		2.014***
		(0.189)		(0.223)
MP			0.138***	0.141***
			(0.044)	(0.044)
IV F-Stat			702.565	698.241
Observations	160,615	160,615	160,615	160,615
Note:				*p<0.1; **p<0.05; ***p<0.01

Table 15: LPM estimates of the non-linear relationship between injury and two week chronic workload with decile and season fixed effects.

	(1)	(2)	(3)	(4)
WHP(3)	0.330^{***}	0.604^{***}	-0.860^{**}	-0.578
	(0.102)	(0.103)	(0.381)	(0.359)
Rested		1.518^{***}		2.005^{***}
		(0.184)		(0.236)
MP			0.130^{***}	0.138^{***}
			(0.040)	(0.041)
IV F-Stat			808.385	790.574
Observations	160,615	160,615	160,615	160,615
Note:				*p<0.1; **p<0.05; ***p<0.01

Table 16: LPM estimates of the linear relationship between injury and three week chronic workload with decile and season fixed effects.

	(1)	(2)	(3)	(4)
WHP(3)	0.092	0.975***	−1.403***	−0.343
	(0.325)	(0.327)	(0.540)	(0.480)
WHP^2(3)	0.122	−0.185	0.274*	−0.116
	(0.162)	(0.162)	(0.165)	(0.161)
Rested		1.566***		2.034***
		(0.185)		(0.236)
MP			0.130***	0.138***
			(0.041)	(0.041)
IV F-Stat			795.583	784.909
Observations	160,615	160,615	160,615	160,615
Note:			*p<0.1; **p<0.05; ***p<0.01	

Table 17: LPM estimates of the non-linear relationship between injury and three week chronic workload with decile and season fixed effects.

	(1)	(2)	(3)	(4)
WHP(4)	0.367^{***}	0.614^{***}	-0.749^{**}	-0.512
	(0.104)	(0.105)	(0.356)	(0.336)
Rested		1.490^{***}		2.035^{***}
		(0.183)		(0.244)
MP			0.123^{***}	0.134^{***}
			(0.038)	(0.039)
IV F-Stat			850.83	830.072
Observations	160,615	160,615	160,615	160,615
Note:				*p<0.1; **p<0.05; ***p<0.01

Table 18: LPM estimates of the linear relationship between injury and four week chronic workload with decile and season fixed effects.

	(1)	(2)	(3)	(4)
WHP(4)	0.173	0.957***	−1.172**	−0.222
	(0.335)	(0.334)	(0.506)	(0.448)
$WHP^2(4)$	0.102	−0.176	0.219	−0.148
	(0.172)	(0.170)	(0.172)	(0.169)
Rested		1.528***		2.065***
		(0.183)		(0.245)
MP			0.124***	0.134***
			(0.038)	(0.039)
IV F-Stat			840.641	827.592
Observations	160,615	160,615	160,615	160,615
Note:				*p<0.1; **p<0.05; ***p<0.01

Table 19: LPM estimates of the non-linear relationship between injury and four week chronic workload with decile and season fixed effects.

	Reason X for extra rest					
	Recovery or Coach		Coach		Exogenous	
	(1)	(2)	(3)	(4)	(5)	(6)
Rested(X)	2.188***	0.926**	0.581	−0.408	−0.645	−0.258
	(0.314)	(0.376)	(0.374)	(0.426)	(0.466)	(1.078)
MP	0.124***	0.108***	0.112***	0.101***	0.108***	0.108***
	(0.033)	(0.032)	(0.034)	(0.034)	(0.031)	(0.031)
MP·Rested(X)		0.082***		0.087***		−0.018
		(0.015)		(0.017)		(0.041)
$\hat{e}$		−0.240***		−0.246***		−0.232***
		(0.034)		(0.035)		(0.032)
IV F-Stat	605.518	605.518	583.497	583.497	607.196	607.196
Observations	160,615	160,615	160,615	160,615	160,615	160,615
Note:						*p<0.1; **p<0.05; ***p<0.01

Table 20: Test for the direction of causality in the relationship between extra rest and player injury risk with decile and season fixed effects.

B.3 LPM-CF injury-workload estimates with player, salary-decile, and season fixed effects

The models in this section are linear probability model (LPM) counterparts to the probit and probit-CF estimates presented in the main text. Except for the specifications in Table 21, all specifications include player, salary-decile, and season fixed effects. All standard errors are clustered at the player-season level and are robust to heteroskedasticity. The stacked generated-regressor adjustment described in Appendix C applies only to the LPM-CF estimates in columns (5)–(6) of Table 21 and columns (2), (4), and (6) of Table 30; these covariance estimates also use the HC1-style multiplier $C/(C-1) \times N/(N-K)$. The remaining columns report clustered OLS or 2SLS standard errors and contain no generated regressor. For readability, all coefficients and standard errors are scaled by 100.

	OLS	OLS	2SLS	2SLS	LPM-CF	LPM-CF
	(1)	(2)	(3)	(4)	(5)	(6)
MP	-0.046^{***}	-0.176^{***}	0.110^{*}	0.105^{***}	0.047	0.072^{**}
	(0.006)	(0.008)	(0.066)	(0.027)	(0.068)	(0.032)
MP^2					0.002^{***}	0.001^{*}
					(0.0004)	(0.0005)
$\hat{e}$					-0.167^{**}	-0.300^{***}
					(0.067)	(0.028)
Constant	4.329^{***}		0.761		1.120	
	(0.153)		(1.510)		(1.514)	
Player, Decile & Season FE		✓		✓		✓
IV F-Stat			69.24	1526.869	69.24	1526.869
Observations	160,615	160,615	160,615	160,615	160,615	160,615
Note:						*p<0.1; **p<0.05; ***p<0.01

Table 21: LPM estimates of risk of injury with respect to time on-court with player, decile and season fixed effects.

	(1)	(2)	(3)	(4)
WHP(1)	0.579***	0.836***	−0.002	0.297
	(0.096)	(0.097)	(0.214)	(0.200)
Rested		1.285***		1.505***
		(0.187)		(0.203)
MP			0.105***	0.105***
			(0.035)	(0.035)
IV F-Stat			1505.233	1508.973
Observations	160,615	160,615	160,615	160,615
Note:				*p<0.1; **p<0.05; ***p<0.01

Table 22: LPM estimates of the linear relationship between injury and one week chronic workload with player, decile and season fixed effects.

	(1)	(2)	(3)	(4)
WHP(1)	−0.056	0.735***	−0.986**	−0.087
	(0.256)	(0.268)	(0.393)	(0.372)
$WHP^2(1)$	0.301***	0.046	0.463***	0.174
	(0.117)	(0.121)	(0.127)	(0.127)
Rested		1.264***		1.428***
		(0.194)		(0.204)
MP			0.106***	0.106***
			(0.035)	(0.035)
IV F-Stat			1491.723	1495.5
Observations	160,615	160,615	160,615	160,615
Note:				*p<0.1; **p<0.05; ***p<0.01

Table 23: LPM estimates of the non-linear relationship between injury and one week chronic workload with player, decile and season fixed effects.

	(1)	(2)	(3)	(4)
WHP(2)	0.757***	0.973***	0.143	0.379*
	(0.107)	(0.107)	(0.229)	(0.214)
Rested		1.205***		1.489***
		(0.185)		(0.210)
MP			0.100***	0.105***
			(0.033)	(0.034)
IV F-Stat			1572.864	1569.805
Observations	160,615	160,615	160,615	160,615
Note:				*p<0.1; **p<0.05; ***p<0.01

Table 24: LPM estimates of the linear relationship between injury and two week chronic workload with player, decile and season fixed effects.

	(1)	(2)	(3)	(4)
WHP(2)	0.116	0.793**	−0.884**	−0.109
	(0.306)	(0.312)	(0.439)	(0.410)
$WHP^2(2)$	0.317**	0.087	0.505***	0.235
	(0.145)	(0.147)	(0.155)	(0.152)
Rested		1.181***		1.425***
		(0.189)		(0.208)
MP			0.101***	0.105***
			(0.034)	(0.034)
IV F-Stat			1564.648	1560.633
Observations	160,615	160,615	160,615	160,615
Note:			*p<0.1; **p<0.05; ***p<0.01	

Table 25: LPM estimates of the non-linear relationship between injury and two week chronic workload with player, decile and season fixed effects.

	(1)	(2)	(3)	(4)
WHP(3)	0.817***	0.994***	0.215	0.396*
	(0.110)	(0.109)	(0.218)	(0.205)
Rested		1.129***		1.471***
		(0.184)		(0.215)
MP			0.098***	0.106***
			(0.032)	(0.032)
IV F-Stat			1631.931	1624.92
Observations	160,615	160,615	160,615	160,615
Note:			*p<0.1; **p<0.05; ***p<0.01	

Table 26: LPM estimates of the linear relationship between injury and three week chronic workload with player, decile and season fixed effects.

	(1)	(2)	(3)	(4)
WHP(3)	0.127	0.664**	−0.805*	−0.178
	(0.325)	(0.327)	(0.426)	(0.398)
$WHP^2(3)$	0.353**	0.166	0.518***	0.288*
	(0.161)	(0.161)	(0.167)	(0.164)
Rested		1.095***		1.414***
		(0.186)		(0.214)
MP			0.100***	0.107***
			(0.032)	(0.032)
IV F-Stat			1623.901	1616.064
Observations	160,615	160,615	160,615	160,615
Note:				*p<0.1; **p<0.05; ***p<0.01

Table 27: LPM estimates of the non-linear relationship between injury and three week chronic workload with player, decile and season fixed effects.

	(1)	(2)	(3)	(4)
WHP(4)	0.902***	1.052***	0.336	0.479**
	(0.111)	(0.110)	(0.204)	(0.194)
Rested		1.091***		1.472***
		(0.183)		(0.220)
MP			0.096***	0.106***
			(0.031)	(0.031)
IV F-Stat			1652.781	1646.122
Observations	160,615	160,615	160,615	160,615
Note:				*p<0.1; **p<0.05; ***p<0.01

Table 28: LPM estimates of the linear relationship between injury and four week chronic workload with player, decile and season fixed effects.

	(1)	(2)	(3)	(4)
WHP(4)	0.128	0.578*	−0.730*	−0.200
	(0.334)	(0.334)	(0.413)	(0.388)
$WHP^2(4)$	0.408**	0.247	0.555***	0.351**
	(0.170)	(0.169)	(0.174)	(0.171)
Rested		1.051***		1.418***
		(0.184)		(0.219)
MP			0.098***	0.107***
			(0.031)	(0.031)
IV F-Stat			1645.226	1637.633
Observations	160,615	160,615	160,615	160,615
Note:			*p<0.1; **p<0.05; ***p<0.01	

Table 29: LPM estimates of the non-linear relationship between injury and four week chronic workload with player, decile and season fixed effects.

	Reason X for extra rest					
	Recovery or Coach		Coach		Exogenous	
	(1)	(2)	(3)	(4)	(5)	(6)
Rested(X)	1.388***	0.463	0.214	−0.936***	−1.205**	−0.389
	(0.235)	(0.317)	(0.241)	(0.314)	(0.468)	(1.086)
MP	0.117***	0.106***	0.107***	0.094***	0.106***	0.107***
	(0.028)	(0.027)	(0.028)	(0.028)	(0.026)	(0.026)
MP·Rested(X)		0.060***		0.101***		−0.038
		(0.016)		(0.017)		(0.042)
$\hat{e}$		−0.308***		−0.306***		−0.301***
		(0.029)		(0.030)		(0.028)
IV F-Stat	1548.405	1548.405	1517.014	1517.014	1537.253	1537.253
Observations	160,615	160,615	160,615	160,615	160,615	160,615
Note:						*p<0.1; **p<0.05; ***p<0.01

Table 30: Test for the direction of causality in the relationship between extra rest and player injury risk with player, decile and season fixed effects.

C Identification and Clustered GMM Inference for the Control-Function Probit

This appendix establishes identification and inference for the probit control-function estimator used in the paper. For inference, we treat the two estimation stages as one exactly identified system of stacked moments. This representation lets us use the clustered laws of large numbers, central limit theorem, and covariance results of Hansen and Lee (2019). This treatment is equivalent to the generated-regressor adjustment of Murphy and Topel (2002) and Wooldridge (2010).

C.1 Model and identification

Index clusters by $c = 1, \ldots, C$ and observations within a cluster by $j = 1, \ldots, n_c$, with $N = \sum_c n_c$. An observation is a player-game and, in the application, a cluster is a player-season. The notation (c, j) therefore indexes the same observations denoted (i, g) in the main text. Let $\text{Injury}_{cj} \in \{0, 1\}$, let MP_{cj} be minutes played, let w_{cj} contain the exogenous controls, and let Z_{cj} be the excluded teammate-absenteeism instrument. Collect the first-stage regressors in the column vector $q_{cj} = (w'_{cj}, Z'_{cj})'$ and write the row vector of second-stage regressors as $x_{cj} = f(w_{cj}, \text{MP}_{cj})$. The model is

$$\text{Injury}_{cj} = \mathbb{1}(x_{cj}\beta + \epsilon_{cj} > 0)\,, \tag{10}$$

$$\text{MP}_{cj} = q'_{cj}\pi + e_{cj}. \tag{11}$$

Assumption C.1 (Control function and rank)**.** The pair (ϵ_{cj}, e_{cj}) is independent of q_{cj} and jointly normal with mean zero, $\text{Var}(\epsilon_{cj}) = 1$, $\text{Var}(e_{cj}) = \sigma_e^2 > 0$, and $|\operatorname{Corr}(\epsilon_{cj}, e_{cj})| < 1$. The excluded-instrument coefficient in π is nonzero. The probability limits

$$Q = \operatorname*{plim}_{N\to\infty} N^{-1} \sum_{c,j} q_{cj} q'_{cj} \quad \text{and} \quad \operatorname*{plim}_{N\to\infty} N^{-1} \sum_{c,j} (x_{cj}, e_{cj})'(x_{cj}, e_{cj})$$

exist and are nonsingular.

The variance normalization fixes the scale of the structural coefficients. The independence condition is stronger than mean independence and is the control-function counterpart of instrument exogeneity used in the main text. The relevance and rank conditions are the analogs of the identification condition in Rivers and Vuong (1988). Because the excluded instrument shifts minutes but does not enter x_{cj}, it provides independent variation in e_{cj}, identifying $\widetilde{\gamma}$ separately from $\widetilde{\beta}$.

Lemma C.2 (Control-function representation)**.** *Let* $\sigma_{\epsilon e} = \mathrm{Cov}(\epsilon_{cj}, e_{cj})$, $\gamma = \sigma_{\epsilon e}/\sigma_e^2$, *and* $\sigma^2 = 1 - \sigma_{\epsilon e}^2/\sigma_e^2$. *Then* $\epsilon_{cj} = \gamma e_{cj} + \xi_{cj}$, *where* $\xi_{cj} \perp\!\!\!\perp (e_{cj}, q_{cj})$ *and* $\xi_{cj} \sim N(0, \sigma^2)$. *Consequently,*

$$\Pr(\mathit{Injury}_{cj} = 1 \mid x_{cj}, e_{cj}) = \Phi(x_{cj}\widetilde{\beta} + e_{cj}\widetilde{\gamma}), \qquad \widetilde{\beta} = \beta/\sigma, \quad \widetilde{\gamma} = \gamma/\sigma. \tag{12}$$

Proof. Set $\xi_{cj} = \epsilon_{cj} - \gamma e_{cj}$. Joint normality and $\mathrm{Cov}(\xi_{cj}, e_{cj}) = 0$ imply $\xi_{cj} \perp\!\!\!\perp e_{cj}$; independence of (ϵ_{cj}, e_{cj}) from q_{cj} then implies $\xi_{cj} \perp\!\!\!\perp (e_{cj}, q_{cj})$. Direct calculation gives $\mathrm{Var}(\xi_{cj}) = 1 - \sigma_{\epsilon e}^2/\sigma_e^2 = \sigma^2$. Because x_{cj} is a function of (q_{cj}, e_{cj}), conditioning equation (10) on (x_{cj}, e_{cj}) and dividing its index by σ yields (12). □

Proposition C.3 (Identification)**.** *Under Assumption C.1,* π, e_{cj}, *and the probit index coefficients* $\theta_0 = (\widetilde{\beta}', \widetilde{\gamma})'$ *are identified. The structural slope* β, *the conditional scale* σ, *and* $\rho = \mathrm{Corr}(\epsilon_{cj}, e_{cj})$ *are separately identified, with*

$$\frac{1}{\sigma} = \sqrt{1 + \sigma_e^2 \widetilde{\gamma}^{\,2}}, \qquad \beta = \sigma\widetilde{\beta} = \frac{\widetilde{\beta}}{\sqrt{1 + \sigma_e^2 \widetilde{\gamma}^{\,2}}}. \tag{13}$$

Proof. Independence gives $\mathrm{E}[q_{cj} e_{cj}] = 0$, so nonsingularity of Q identifies the population least-squares coefficient π and hence the residual $e_{cj} = \mathrm{MP}_{cj} - q_{cj}'\pi$; its variance $\sigma_e^2 = \mathrm{E}[e_{cj}^2]$ is then identified as the second moment of that residual. If two values of θ generated the same conditional probability in (12), strict monotonicity of Φ would make their indices equal almost surely. The pooled second moment of (x_{cj}, e_{cj}) is the probability limit that averages

over the heterogeneous within-cluster distributions. Its nonsingularity under Assumption C.1 then makes the coefficients equal. This is the residual-inclusion identification argument of Rivers and Vuong (1988).

It remains to undo the conditional scale normalization. Since $\gamma = \rho/\sigma_e$ and $\sigma = \sqrt{1-\rho^2}$,

$$\sigma_e \widetilde{\gamma} = \frac{\rho}{\sqrt{1-\rho^2}}.$$

Thus the magnitude and sign of ρ are identified from σ_e and $\widetilde{\gamma}$, and solving for σ gives (13). This is the scalar control-function scale recovery discussed by Wooldridge (2015). □

Here and in the main text, σ is the standard deviation of the conditional error ξ_{cj}, not the standard deviation of ϵ_{cj}, which is normalized to one. Thus the reported $1/\sigma$ and the structural coefficient β retain the same meaning across specifications.

C.2 Consistency, asymptotic normality, and inference

Let $d_{cj} = (\text{Injury}_{cj}, \text{MP}_{cj}, q_{cj})$ and $e_{cj}(\pi) = \text{MP}_{cj} - q_{cj}'\pi$. Let $s(d_{cj}, \theta; \pi)$ be the column-vector probit score in θ, evaluated at $x_{cj}\widetilde{\beta} + e_{cj}(\pi)\widetilde{\gamma}$. Define

$$\alpha = (\pi', \theta')', \qquad m_{cj}(\alpha) = \begin{pmatrix} q_{cj} e_{cj}(\pi) \\ s(d_{cj}, \theta; \pi) \end{pmatrix}, \qquad \bar{m}_N(\alpha) = \frac{1}{N}\sum_{c,j} m_{cj}(\alpha).$$

The number of moments equals the dimension of α. The first block of $\bar{m}_N(\widehat{\alpha}) = 0$ is the standard OLS first-order condition, and the second block is the probit score equation using $\widehat{e}_{cj} = e_{cj}(\widehat{\pi})$. When the probit log-likelihood has an interior maximizer, its first-order condition is the second block above. The sequential OLS/probit estimates therefore solve $\bar{m}_N(\widehat{\alpha}) = 0$, allowing us to study their asymptotic properties through this stacked, exactly identified system.

Write

$$H = \operatorname{plim} \frac{1}{N} \sum_{c,j} \left(-\frac{\partial s(d_{cj}, \theta_0; \pi)}{\partial \theta'} \right), \qquad F = \operatorname{plim} \frac{1}{N} \sum_{c,j} \frac{\partial s(d_{cj}, \theta_0; \pi)}{\partial \pi'}. \tag{14}$$

The Jacobian of the stacked moments at the true parameters α_0 is block triangular:

$$A = \operatorname{plim} \frac{\partial \bar{m}_N(\alpha_0)}{\partial \alpha'} = \begin{pmatrix} -Q & 0 \\ F & -H \end{pmatrix}. \tag{15}$$

The signs follow from $\partial(q_{cj} e_{cj})/\partial \pi' = -q_{cj} q_{cj}'$ and the definition of H.

Assumption C.4 (Clustered GMM regularity)**.** Clusters are independent, while dependence within a cluster is unrestricted. Cluster sizes may differ but are uniformly bounded; hence $\max_c n_c/N \to 0$, $\sum_c n_c^2/N = O(1)$, and $\max_c n_c^2/N \to 0$. The parameter space is compact, α_0 is an interior and uniquely identified root, and A is nonsingular. The sample probit log-likelihood has an interior maximizer with probability approaching one. The stacked moments and their derivatives satisfy the continuity, envelope, moment, uniform integrability, and eigenvalue conditions of Hansen and Lee (2019), Theorems 12 and 13. The cluster-moment covariance and the resulting sandwich covariance defined below are uniformly nonsingular with bounded spectrum: their eigenvalues lie in a fixed positive interval for all N. These maintained bounds are used for the inference below.

Bounded cluster size is natural here because a player-season contains at most one season's games. It is stronger than necessary, but gives a transparent sufficient condition for the unequal-cluster-size restrictions while allowing arbitrary dependence among observations in the same player-season.

For cluster c, set

$$M_c(\alpha) = \sum_{j=1}^{n_c} m_{cj}(\alpha).$$

Define the finite-N stacked-moment covariance:

$$B_N = \frac{1}{N}\sum_c \mathrm{E}[(M_c(\alpha_0) - \mathrm{E}\, M_c(\alpha_0))(M_c(\alpha_0) - \mathrm{E}\, M_c(\alpha_0))'] \,. \tag{16}$$

Note that B_N need not converge to a fixed limit.

Theorem C.5 (Consistency, asymptotic normality, and inference)**.** *Under Assumptions C.1 and C.4,* $\widehat{\alpha} \xrightarrow{p} \alpha_0$*, and with the stacked-GMM sandwich* $V_N = A^{-1}B_N A^{-1\prime}$*,*

$$V_N^{-1/2}\sqrt{N}(\widehat{\alpha} - \alpha_0) \xrightarrow{d} N(0, I). \tag{17}$$

The sandwich is estimated by its sample analog: with

$$\widehat{A} = \frac{1}{N}\sum_{c,j} \frac{\partial m_{cj}(\widehat{\alpha})}{\partial \alpha'}, \tag{18}$$

$$\widehat{B} = \frac{1}{N}\sum_c M_c(\widehat{\alpha})M_c(\widehat{\alpha})' - \frac{1}{N}\sum_c n_c^2 \bar{m}_N(\widehat{\alpha})\bar{m}_N(\widehat{\alpha})', \tag{19}$$

$$\widehat{V} = \widehat{A}^{-1}\widehat{B}\widehat{A}^{-1\prime}, \tag{20}$$

More generally, for any fixed full-rank $k \times q$ *contrast matrix* R*,*

$$(R'V_N R)^{-1/2} R'\widehat{V} R (R'V_N R)^{-1/2} \xrightarrow{p} I_q, \tag{21}$$

$$(R'V_N R)^{-1/2} R'\sqrt{N}(\widehat{\alpha} - \alpha_0) \xrightarrow{d} N(0, I_q), \qquad (R'\widehat{V}R)^{-1/2} R'\sqrt{N}(\widehat{\alpha} - \alpha_0) \xrightarrow{d} N(0, I_q), \tag{22}$$

Equation (17) *is the identity-contrast case* $R = I$*. Thus* $N^{-1}\widehat{V}$ *is the covariance estimator for* $\widehat{\alpha}$*, feasible t-statistics for scalar contrasts* ($q = 1$) *are asymptotically standard normal, and fixed-dimensional Wald statistics are asymptotically chi-squared by the continuous mapping theorem.*

Proof. We adapt the clustered GMM proofs of Hansen and Lee (2019), matching their parameter with α, their sample size with N, and their weighting matrix with the identity matrix. Assumption C.4 satisfies their maintained conditions for this problem, so their

proof of Theorem 12 gives $\widehat{\alpha} \xrightarrow{p} \alpha_0$.

Their proof of Theorem 13 supplies the limit distribution. Two objects in our theorem's statement must be matched to theirs. First (using Hansen and Lee's notation, with N in place of their n), their Ω_N, the uncentered covariance of the cluster moments, equals our B_N because $\mathrm{E}[M_c(\alpha_0)] = 0$ makes the centered form (16) coincide with the uncentered one; the mean is zero term by term, since the OLS moment $q_{cj}e_{cj}$ is orthogonal by construction and the probit score is mean zero at θ_0 by Lemma C.2.[20]

Second (continuing in Hansen and Lee's notation), the finite-N expected Jacobian is $Q_N = \mathrm{E}[\partial \bar{m}_N(\alpha_0)/\partial \alpha']$; its sample analog $\widehat{Q}_N(\widehat{\alpha})$ is our $\widehat{A}$. The Jacobian law-of-large-numbers and uniform-integrability conditions in Assumption C.4 imply $Q_N \to A$. For our exactly identified system, Hansen and Lee's covariance based on Q_N reduces to $V_N^{HL} = Q_N^{-1} B_N Q_N^{-1\prime}$. Since $Q_N \to A$, A is nonsingular, and the relevant spectra are uniformly bounded, $\|V_N^{HL} - V_N\| \to 0$, where $V_N = A^{-1} B_N A^{-1\prime}$. The corresponding fixed-contrast standardizations in the statement of our theorem are therefore asymptotically equivalent to those in Hansen and Lee (2019). The derivations underlying their (41)-(43) allow any full-rank contrast R_n; applying them at the identity contrast gives (17) and its feasible form, and at a general contrast gives (21) and (22).

Hansen and Lee state Theorems 12 and 13 for the overidentified case $l > k$, whereas our stacked system is exactly identified, $l = k$. Adapting their proof to our case is immediate and in fact removes steps rather than adding them. A fixed positive-definite weighting matrix enters their argument only through the curvature of the GMM criterion; under exact identification the estimator is the weight-free root of $\bar{m}_N(\widehat{\alpha}) = 0$, so the weight cancels from the estimator and, in the sandwich $A^{-1} B_N A^{-1\prime}$, from its covariance. The only place their proof uses $l > k$ is the rank-$(l-k)$ overidentification projection behind the J statistic, which is irrelevant when $l = k$. □

[20] Note that the same remark applies to $\widehat{B}$, the feasible analog of B_N. Equation (19) is Hansen and Lee's centered clustered covariance estimator; the subtracted term is their recentering, which matters only away from the root. Because our estimator is the exact root, $\bar{m}_N(\widehat{\alpha}) = 0$ in sample, that term is zero, and $\widehat{B}$ reduces to the ordinary outer product of the full stacked cluster moments.

C.3 Inference on the probit coefficients

The object of interest is the probit index θ, not the first-stage nuisance π. The next corollary extracts the lower block of Theorem C.5 for $\widehat{\theta}$ alone, and shows that its generated-regressor covariance is already contained in the stacked sandwich, with no separate two-step calculation.

Corollary C.6 (Inference on the probit coefficients)**.** *Under Assumptions C.1 and C.4,*

$$V_{N\theta}^{-1/2}\sqrt{N}(\widehat{\theta}-\theta_0) \xrightarrow{d} N(0,I), \qquad V_{N\theta} = H^{-1}\left\{\frac{1}{N}\sum_c \mathrm{Var}(h_c)\right\}H^{-1\prime}, \tag{23}$$

where $h_c = \sum_j [s_{cj} + FQ^{-1}q_{cj}e_{cj}]$ sums the score $s_{cj} = s(d_{cj},\theta_0;\pi)$ and its first-stage correction within a cluster. The matrix $V_{N\theta}$ is the $\theta\theta$ block of V_N, consistently estimated by the feasible sandwich built from the sample analogs of (14)*,*

$$\widehat{H} = \frac{1}{N}\sum_{c,j}\left(-\frac{\partial s(d_{cj},\widehat{\theta};\widehat{\pi})}{\partial\theta'}\right), \qquad \widehat{F} = \frac{1}{N}\sum_{c,j}\frac{\partial s(d_{cj},\widehat{\theta};\widehat{\pi})}{\partial\pi'}, \qquad \widehat{Q} = \frac{1}{N}\sum_{c,j} q_{cj}q_{cj}', \tag{24}$$

namely

$$\widehat{V}_\theta = \widehat{H}^{-1}\left\{\frac{1}{N}\sum_c \widehat{h}_c\widehat{h}_c'\right\}\widehat{H}^{-1\prime}, \qquad \widehat{h}_c = \sum_j\left[\widehat{s}_{cj} + \widehat{F}\widehat{Q}^{-1}q_{cj}\widehat{e}_{cj}\right], \tag{25}$$

with fitted score $\widehat{s}_{cj} = s(d_{cj},\widehat{\theta};\widehat{\pi})$ and first-stage residual $\widehat{e}_{cj} = e_{cj}(\widehat{\pi})$. This $\widehat{V}_\theta$ is exactly the $\theta\theta$ block of $\widehat{V}$ in (20)*; thus $N^{-1}\widehat{V}_\theta$ is the covariance estimator for $\widehat{\theta}$*[21]*, feasible t-statistics for scalar contrasts in θ are asymptotically standard normal, and fixed-dimensional Wald statistics in θ are asymptotically chi-squared.*

Proof. Block inversion of (15) gives

$$A^{-1} = \begin{pmatrix} -Q^{-1} & 0 \\ -H^{-1}FQ^{-1} & -H^{-1}\end{pmatrix},$$

[21]In the empirical implementation, this covariance is multiplied by $C/(C-1)\times N/(N-K)$, where C is the number of player-season clusters and $K = \dim(\theta)$. This HC1-style finite-sample adjustment converges to one under the maintained asymptotics and therefore leaves the estimator above unchanged.

so the θ block of the influence function $-A^{-1}m_{cj}(\alpha_0)$ is

$$H^{-1}\Big[s_{cj} + FQ^{-1}q_{cj}e_{cj}\Big]. \tag{26}$$

Summing within a cluster gives h_c, so the $\theta\theta$ block of $V_N = A^{-1}B_N A^{-1\prime}$ is

$$H^{-1}\left\{\frac{1}{N}\sum_c \mathrm{Var}(h_c)\right\}H^{-1\prime},$$

which is $V_{N\theta}$ in (23).

The sample Jacobian $\widehat{A}$ in (18) has the same block-triangular pattern, $\widehat{A} = \begin{pmatrix} -\widehat{Q} & 0 \\ \widehat{F} & -\widehat{H} \end{pmatrix}$ with $\widehat{H}, \widehat{F}, \widehat{Q}$ as in (24), because differentiating $m_{cj}(\alpha) = (q_{cj}e_{cj}(\pi), s(d_{cj}, \theta; \pi))'$ and evaluating at $\widehat{\alpha}$ reproduces (15) with (θ_0, π) replaced by $(\widehat{\theta}, \widehat{\pi})$. The identical block inversion therefore gives $\widehat{A}^{-1} = \begin{pmatrix} -\widehat{Q}^{-1} & 0 \\ -\widehat{H}^{-1}\widehat{F}\widehat{Q}^{-1} & -\widehat{H}^{-1} \end{pmatrix}$, so the θ block of $-\widehat{A}^{-1}m_{cj}(\widehat{\alpha})$ is $\widehat{H}^{-1}[\widehat{s}_{cj} + \widehat{F}\widehat{Q}^{-1}q_{cj}\widehat{e}_{cj}]$, which sums within a cluster to $\widehat{h}_c$. Because our estimator is the exact root, the recentering term in (19) is zero, so $\widehat{B} = \frac{1}{N}\sum_c M_c(\widehat{\alpha})M_c(\widehat{\alpha})'$, and the same sandwich algebra gives the $\theta\theta$ block of $\widehat{V} = \widehat{A}^{-1}\widehat{B}\widehat{A}^{-1\prime}$ as $\widehat{H}^{-1}\{\frac{1}{N}\sum_c \widehat{h}_c\widehat{h}_c'\}\widehat{H}^{-1\prime}$, which is $\widehat{V}_\theta$ in (25).

The standardized law and the consistency of $\widehat{V}_\theta$ are the $\theta\theta$ restrictions of Theorem C.5, and the bounded spectrum of V_N passes to its principal block $V_{N\theta}$. Extracting θ from α is itself the contrast $R = S$ for the fixed selection matrix S with $\widehat{\theta} = S'\widehat{\alpha}$, so (22) at $R = S$ delivers the feasible conclusions here. □

C.4 Predicted probabilities and marginal effects

The paper reports predicted probabilities and marginal effects evaluated at selected covariate points. For the inference result below, treat the evaluation point (x^*, e^*) as fixed. Let $a = x^*\widetilde{\beta} + e^*\widetilde{\gamma}$, and let $\mu(\theta)$ denote the resulting smooth scalar effect. For a predicted probability,

$$\mu(\theta) = \Phi(a), \qquad \nabla\mu(\theta) = \phi(a)(x^*, e^*)'.$$

For a continuous regressor entering a linearly, order the coordinates of $\widetilde{\beta}$ and x^*, without loss of generality, so that the regressor enters first, with coefficient $\widetilde{\beta}_1$,

$$\mu(\theta) = \phi(a)\widetilde{\beta}_1, \qquad \frac{\partial \mu}{\partial \widetilde{\beta}_\ell} = \phi(a)\mathbb{1}\{\ell = 1\} - a\phi(a)\widetilde{\beta}_1 x^*_\ell, \qquad \frac{\partial \mu}{\partial \widetilde{\gamma}} = -a\phi(a)\widetilde{\beta}_1 e^*,$$

with ℓ indexing the coordinates of $\widetilde{\beta}$. When the regressor and its square are both included as the first and the second components in x, the derivatives are

$$\begin{aligned}
\frac{\partial \mu}{\partial \widetilde{\beta}_\ell} &= \phi(a)\Big(\mathbb{1}\{\ell = 1\} + 2x^*_1\mathbb{1}\{\ell = 2\}\Big) - a\phi(a)\Big(\widetilde{\beta}_1 + 2\widetilde{\beta}_2 x^*_1\Big)x^*_\ell, \\
\frac{\partial \mu}{\partial \widetilde{\gamma}} &= -a\phi(a)\Big(\widetilde{\beta}_1 + 2\widetilde{\beta}_2 x^*_1\Big)e^*.
\end{aligned}$$

Proposition C.7 (Inference on smooth scalar effects)**.** *Under Assumptions C.1 and C.4, let $\mu(\theta)$ be continuously differentiable in a neighborhood of θ_0 with $\nabla\mu(\theta_0) \neq 0$, and let $\widehat{V}_\theta$ be the $\theta\theta$ block of $\widehat{V}$, as in Corollary C.6. With*

$$\widehat{\mathrm{se}}\{\mu(\widehat{\theta})\} = \left\{\frac{1}{N}\nabla\mu(\widehat{\theta})'\widehat{V}_\theta\nabla\mu(\widehat{\theta})\right\}^{1/2}, \tag{27}$$

the statistic satisfies

$$\frac{\mu(\widehat{\theta}) - \mu(\theta_0)}{\widehat{\mathrm{se}}\{\mu(\widehat{\theta})\}} \xrightarrow{d} N(0,1).$$

Proof. Because μ is nonlinear and $\widehat{\mathrm{se}}\{\mu(\widehat{\theta})\}$ uses the estimated gradient $\nabla\mu(\widehat{\theta})$, this is not immediate from Theorem C.5, whose contrasts are nonrandom, so the proof proceeds in four steps.

First, let S be the fixed matrix that selects θ from α, such that $\widehat{\theta} = S'\widehat{\alpha}$, and set $R = S\nabla\mu(\theta_0)$. Since $R'V_N R = \nabla\mu(\theta_0)'V_{N\theta}\nabla\mu(\theta_0)$, where $V_{N\theta} = S'V_N S$ is the $\theta\theta$ block of V_N by construction of S, and $R'\sqrt{N}(\widehat{\alpha} - \alpha_0) = \nabla\mu(\theta_0)'\sqrt{N}(\widehat{\theta} - \theta_0)$, substituting into (22) gives

$$\left\{\nabla\mu(\theta_0)'V_{N\theta}\nabla\mu(\theta_0)\right\}^{-1/2}\nabla\mu(\theta_0)'\sqrt{N}(\widehat{\theta} - \theta_0) \xrightarrow{d} N(0,1).$$

With ω_N^2 labeling the bracketed term, this becomes $\omega_N^{-1}\nabla\mu(\theta_0)'\sqrt{N}(\widehat{\theta} - \theta_0) \xrightarrow{d} N(0,1)$; ω_N is

bounded away from zero and infinity because $V_{N\theta}$ has bounded, uniformly positive spectrum by Corollary C.6 and because $\nabla\mu(\theta_0) \neq 0$ is assumed.

Second, smoothness of μ and $\sqrt{N}(\widehat{\theta} - \theta_0) = O_p(1)$—which the standardized law (23) delivers together with the bounded norm of $V_{N\theta}$—give the expansion $\sqrt{N}\{\mu(\widehat{\theta}) - \mu(\theta_0)\} = \nabla\mu(\theta_0)'\sqrt{N}(\widehat{\theta}-\theta_0)+o_p(1)$. Third, (21) at $R = S$ gives $V_{N\theta}^{-1/2}\widehat{V}_\theta V_{N\theta}^{-1/2} \xrightarrow{p} I$, so $\widehat{V}_\theta$ has bounded norm in probability and $\widehat{V}_\theta - V_{N\theta} \xrightarrow{p} 0$ (using $V_{N\theta}$'s bounded spectrum from Corollary C.6). Fourth, $\nabla\mu(\widehat{\theta}) \xrightarrow{p} \nabla\mu(\theta_0)$ by continuity.

By (27), $N\,\widehat{\mathrm{se}}\{\mu(\widehat{\theta})\}^2 = \nabla\mu(\widehat{\theta})'\widehat{V}_\theta\nabla\mu(\widehat{\theta})$. Steps 3 and 4 control its difference from ω_N^2 directly, without requiring $V_{N\theta}$ to have a fixed limit: adding and subtracting $\nabla\mu(\theta_0)'\widehat{V}_\theta\nabla\mu(\theta_0)$,

$$N\,\widehat{\mathrm{se}}\{\mu(\widehat{\theta})\}^2 - \omega_N^2 = \underbrace{[\nabla\mu(\widehat{\theta}) - \nabla\mu(\theta_0)]'\widehat{V}_\theta[\nabla\mu(\widehat{\theta}) + \nabla\mu(\theta_0)]}_{(i)} + \underbrace{\nabla\mu(\theta_0)'[\widehat{V}_\theta - V_{N\theta}]\nabla\mu(\theta_0)}_{(ii)}.$$

Term (i) is $o_p(1)$: step 4 gives a vanishing first factor, while $\widehat{V}_\theta$ and $\nabla\mu(\widehat{\theta}) + \nabla\mu(\theta_0)$ stay bounded in probability by step 3 and continuity. Term (ii) is $o_p(1)$: step 3's vanishing matrix difference is pre- and post-multiplied by the fixed vector $\nabla\mu(\theta_0)$. Since ω_N is bounded away from zero, $N\,\widehat{\mathrm{se}}\{\mu(\widehat{\theta})\}^2/\omega_N^2 \xrightarrow{p} 1$.

Finally, we apply Slutsky's theorem twice to assemble these results into the final statement of the proposition. First, adding the second step's $o_p(1)$ remainder to the first step's limit does not change it, since ω_N is bounded away from zero. Thus $\sqrt{N}\{\mu(\widehat{\theta})-\mu(\theta_0)\}/\omega_N \xrightarrow{d} N(0,1)$. Second, the exact factorization

$$\frac{\mu(\widehat{\theta}) - \mu(\theta_0)}{\widehat{\mathrm{se}}\{\mu(\widehat{\theta})\}} = \frac{\sqrt{N}\{\mu(\widehat{\theta}) - \mu(\theta_0)\}}{\omega_N}\left\{\frac{N\,\widehat{\mathrm{se}}\{\mu(\widehat{\theta})\}^2}{\omega_N^2}\right\}^{-1/2}$$

multiplies this statistic by the reciprocal square root of the ratio derived above. That factor converges in probability to one, so a second application of Slutsky's theorem gives the claimed standard normal limit. □