

Playing the Spread: The Implications of Worker Volatility for Human Capital Management

Isaac Tucker

Abstract

Workers have good days and bad days, but human capital management research overwhelmingly focuses on mean performance. This paper examines whether organizations recognize and price worker volatility when they can measure it, and whether they should. In two studies using data from professional basketball, I document both wage and normative penalties for volatile workers. After separating persistent volatility from playing-time noise, more volatile players earn less on the margin where contracts are negotiable. A difference-in-differences design then exploits the quasi-random timing of absences as an identifying shock. Holding mean performance fixed, losing a reliable player inflicts a larger penalty on team point differential and win probability than losing a volatile one. I show that the steadiest players are nearly as costly to lose as high-mean stars, at forty percent of the salary. Taken together, these results suggest that performance management is a portfolio problem and worker volatility is an important locus for strategic action.

1. Introduction

On June 11, 2013, the Miami Heat lost in Game 3 of the National Basketball Association (NBA) finals by 36 points. LeBron James – the league’s most valuable player, in the midst of perhaps the greatest individual season in basketball history, leading a super-team with the NBA’s best record – scored 15 points on 21 shot attempts. It was his worst performance of the season. Advanced statistics suggest he played at a level undeserving of a roster slot (Myers 2020). With the benefit of hindsight, we know that this was an aberration: the Heat would go on to win the best-of-seven championship series in Game 7 behind 37 points from their superstar. But, in all likelihood, there was at least one person among Game 3’s 14 million television viewers who had never watched professional basketball before. A neophyte, lacking context, might reasonably have concluded that one of the greatest players of all time was a mediocre talent on an over-matched team.

Though absurd in the context of professional basketball, the measurement and management of stochastic performance is a material and under-addressed challenge. Organizations routinely evaluate workers on singular or sparse events, despite evidence that these are noisy signals with limited predictive value. Hiring, retention, termination, promotion, and performance bonus decisions respond to periods of recent over-performance and punish underachievers. And yet, bad days are all too common. Whether because the baby won’t sleep, the weather won’t cooperate, or the muse won’t sing, individual performance is subject to nadirs at varying scales and intervals. Meta-analyses suggest that roughly half of variation in performance outcomes in both lab and field settings is attributable to within-person volatility, rather than between-person differences (Dalal et al. 2020; Podsakoff et al. 2019). In the case of professional basketball, 76–92% of total game-over-game variation is within player-season.¹

Since 2000, six top management journals² have published 316 empirical papers containing any mention of human resources or human capital management alongside any mention of per-

¹Specifically, this is variation in performance among NBA starters (639 players, 143,064 player-games, 2011–2023) for rate-adjusted Game Score (76%), BPM (86%), or plus/minus (92%).

²AMJ, AMR, ASQ, Management Science, Organization Science, SMJ

formance or productivity³. This corpus contains a wealth of measures and methods, but not a single paper using individual, within-person, intertemporal variance as a focal outcome or treatment variable. This is not to say that the role of volatility is unknown. An adjacent tradition in industrial and organizational psychology treats performance variability as a distinct dimension of worker performance (Sackett et al. 1988) and finds that it interacts with compensation (Awtrey et al. 2021; Barnes & Morgeson 2007). Some authors acknowledge the absence of volatility in their analysis (De Figueiredo et al. 2013). Others include it in their theoretical specification but not their empirical findings (Aral et al. 2012; Battiston et al. 2025; DeVaro 2006).⁴ It is possible that this is a costless omission. However, stochastic performance can impose a non-negligible (Lazear 2004) and sometimes catastrophic (Perrow 1984) cost on the organization. The extent to which these costs impact organizational performance and, therefore, strategy is largely unexplored.

The omission is particularly important for strategic human capital management. Unlike other assets, human capital is embodied in workers who firms do not own and who can leave of their own volition or bargain over the value they create (Coff 1997, 1999). Decision-makers must identify which worker features create value for the organization and determine how those characteristics are valued on the labor market. Where there is a gap in the private and public valuations of a worker, the organization can extract value and build a sustainable competitive advantage. Traditionally, this is the result of bundles of HR practices or the development of firm-specific human capital (Becker & Gerhart 1996; Bidwell 2011; Chadwick & Dabu 2009). If the spread of a worker's performance has value distinct from the mean of their performance, it would represent an additional (and novel) wedge between wages and performance.

This paper asks two questions: do organizations price worker volatility when they are able to measure it? And does it matter for organizational performance? The answers dictate how we interpret the absence, in both research and practice, of a coherent human capital management

³On Web of Science, the query is (human resources OR human capital) AND (productivity OR performance)

⁴Only a handful of researchers concerned with the effectiveness of healthcare operations have assessed the effect of policies, interventions, and experience on worker volatility (Bavafa & Jónasson 2021, 2024; Lu et al. 2022). However, because this work is addressed to the operations management literature, it does not show up in the initial literature search and may not be readily apparent to HR researchers or practitioners.

narrative around worker volatility. If steadiness and reliability are valuable but unvalued, this suggests a bias or blind-spot among decision-makers. If, instead, a worker's bad days have no bearing on the organization as a whole, then worker volatility can be safely disregarded as a locus for strategic action.

I answer these questions using data from the National Basketball Association. Study 1 shows that teams price consistency among players whose compensation is negotiated rather than fixed by the league's minimum and maximum salary rules. Study 2 exploits the quasi-random timing of player absences to estimate the corresponding production value of reliability. Holding mean performance fixed, losing a more reliable rotation player imposes a larger penalty on team point differential and win probability than losing a more volatile one. The implied wage and production tradeoffs are consistent, suggesting that teams price reliability because it contributes to organizational performance.

This paper makes three contributions to research on strategic human capital. First, it provides evidence that within-person performance volatility has consequences for both compensation and organizational performance. Second, I show that organizations use information (when it is available) about the distribution of worker performance when making strategic human capital decisions. Finally, the results help explain why volatility has received so little attention in prior research. Even in an unusually data-rich setting, reliability is difficult to distinguish from playing-time noise and its wage consequences emerge only over longer evaluation windows. The absence of volatility from existing HCM research may therefore reflect the difficulty of measuring it more than its lack of strategic importance.

2. Research Setting

Professional basketball has a long history as an empirical setting for questions about performance and randomness. Since Gilovich et al.'s (1985) foundational paper on shooting sequences, a large literature in psychology and sports analytics has used NBA data to study whether player output is more or less random than intuition suggests. Unlike that research, which uses the possession

(really, the shot) as the unit of analysis, I examine performance at the level of a single game. One game gives a player ample opportunity to impact the outcome whereas a single possession can occur without the involvement of 9 out of 10 players in the game. Game performance has an advantage over season-level performance in that every game has a winner, a loser, and a common set of environmental and temporal conditions under which it is played. Still, this is a choice made (in part) for convenience. Volatility occurs from moment to moment (Dalal et al. 2020) and every period is the convolution of these micro-events. As our understanding of performance grows, it should naturally direct more attention to the scale at which workers are measured.

2.1. Measuring Performance in Professional Basketball

Professional basketball teams play 82-game seasons. Each game lasts 48 minutes and the court is shared by five players on each team, giving 240 team-minutes per game. Teams take turns attempting to score baskets. An average game sees just over 100 possessions per team. On each possession, players can contribute to the success or failure of their team's attempt to score in multiple ways. Many of these are captured and tracked with basic counting statistics. In addition to points scored, players are measured on their shooting percentage (shots made divided by shots attempted), assists (passes which lead to teammate scores), blocks and steals (defensive actions which prevent shooting opportunities for the other team), turnovers (possession lost without a shot attempt) and rebounds (ball recoveries after a missed shot).

As in many other jobs, basketball players have different strengths and inclinations, which frustrates attempts to collapse these statistics into a single number. Point guards, responsible for distributing the ball to others and chosen for their passing ability, will naturally get more assists than players who are on the court because they are tall enough to grab rebounds. The simplest approach, and the one I use in this paper, is a weighted sum of the player's activities called Game Score (Hollinger 2003). The idea behind Game Score is that each possession generates, on average, about one point. So an action which earns the team another possession is worth about a point while one that loses possession subtracts a point. These fractional "points" are then added to the player's

point total, which gives their Game Score. The full formula is:

$$\begin{aligned}\text{Game Score} = & \text{Points Scored} + (0.4 \times \text{Field Goals Made}) - (0.7 \times \text{Field Goal Attempts}) \\ & - (0.4 \times (\text{Free Throw Attempts} - \text{Free Throws})) + (0.7 \times \text{Offensive Rebounds}) \\ & + (0.3 \times \text{Defensive Rebounds}) + \text{Steals} + (0.7 \times \text{Assists}) + (0.7 \times \text{Blocks}) \\ & - (0.4 \times \text{Personal Fouls}) - \text{Turnovers}.\end{aligned}$$

Game Score is not a perfect summary performance measure. It ignores context (opponent quality, pace of play, location of game, quality of teammates, etc.). However, for the study of game-level performance variance it has several important features. Game Score is universally available, easily interpretable from game to game and has obvious face value. It penalizes bad things, rewards good things, and captures the many different ways to add value as a professional basketball player. A player whose Game Score falls from 30 to 5 in back-to-back games had one good game and one bad game, regardless of how they accumulated those statistics.

In order to ensure that the pieces of performance that are not measured in Game Score do not materially alter the below results, I also make use of one additional “advanced” statistic in the form of Box Plus Minus (BPM). This regression-based Game Score analog takes its name from plus/minus, the team’s point differential while a player is on the floor, and estimates how much a player contributes to the team’s margin of victory, on average, across an entire season (Myers 2020). It includes much of the context that Game Score omits but is not necessarily interpretable at the game level. Where it is necessary to control for long-run player quality, I use BPM. Otherwise, I use BPM as a robustness check for Game Score based results.

The other problem with Game Score is that it is a counting statistic, so game-over-game variation is a function of both player characteristics and minutes played. This is akin to binomial noise. Consider, for example, that a single shot attempt has an average success rate of 46%. A player who takes 15 shots per game has a per-game sampling error of roughly 13 percentage points. At the league average of 9 shot attempts per game, this grows to roughly 16 percentage points

(nearly three times the between-player standard deviation in shot percentage after correcting for this binomial noise). Prior work using NBA data (Awtrey et al. 2021; Barnes & Morgeson 2007) makes no correction or adjustment for this noise and so misspecifies the volatility of player performance.

Subsequent analyses use a minutes-adjusted Game Score standard deviation as an estimate of player-specific variance, net of the mechanical variance associated with limited playing time. For the median player-season (23 minutes per game), this correction removes roughly half of the observed variance. I first normalize Game Score per 36 minutes, which is conventional for NBA statistics and reflects the typical regular-season workload for a top player. I then decompose the variance of per-36 Game Score across games within a season into a player-specific characteristic and accompanying sampling noise:

$$\text{Var}_{\text{obs},i}(GS) = \underbrace{\sigma_{\text{true},i}^2}_{\text{player volatility}} + \underbrace{k \cdot \mathbb{E}[1/m_{i,g}]}_{\text{noise}}. \quad (1)$$

Here, m_g is minutes in game g , and $k \cdot \mathbb{E}[1/m_g]$ is finite-exposure noise that shrinks as playing time grows. I estimate k by regressing observed variance on the player’s average inverse minutes for player-seasons in which a player plays more than 8 minutes in more than 20 games. The slope, $\hat{k}$, is approximately 650 (SE 21.7).⁵ Each player’s true variance is then his observed variance less this noise:

$$\hat{\sigma}_{\text{true},i}^2 = \max\{ \text{Var}_{\text{obs},i} - \hat{k} \mathbb{E}_i[1/m_{i,g}], 0 \}. \quad (2)$$

Hereafter, a player’s volatility refers to the square root of this adjusted variance.

Figure 1 demonstrates the importance of this correction. In the raw data, players who play less often look more volatile than those who play more minutes. If we had a prior belief that volatility was bad and steadier players should play more often, this would offer confirmation of that

⁵The slope $\hat{k}$ is identified from cross-player variation and inherits any correlation between volatility and playing time. Within-player and composition-based estimates of the same slope are larger and so would remove some volatility which persists within player across seasons. The appendix details this procedure in full and provides a split sample estimate to gauge the magnitude of measurement error.




Figure 1: Observed and adjusted standard deviation of per-36 Game Score, by minutes per game. Points are bin means. The shaded region is the mechanical component removed by Equation (2). Chart shows player-seasons with at least twenty games of eight or more minutes, 2011–2023 ($n = 4,903$).

bias. After correcting for finite-sample noise, the relationship is roughly flat. If anything, the slight uptick in adjusted volatility suggests that the most important players are more volatile than those who are played sparingly.

Roughly 5% of player-seasons have observed variance below the estimated noise, given their minutes played. These players are so consistent that their game-over-game variation in performance is indistinguishable from the “bouncing ball” luck inherent to basketball. Figure 2 shows the spread of player-seasons from the sample period.

2.2. Compensation and the Collective Bargaining Agreement

In the NBA, the rules governing on-court play and roster-building are negotiated between the league (representing the owners of the 30 franchise teams) and representatives for the players’ union (National Basketball Players Association or NBPA) as part of a collective bargaining agreement (CBA). Each CBA is an opportunity to substantially reshape the league landscape and the




Figure 2: Scatter of all qualifying player-seasons, 2011-2023 ($n = 3,724$) by mean per-36 Game Score and adjusted volatility. Labeled points from the 2017-2018 Golden State Warriors.

relationship between workers, fans, and teams. As such, the league’s history is typically broken up into eras defined by the prevailing CBA. This analysis is concerned with the “modern” NBA era, which began with the 2011 CBA and was extended through the 2023-2024 season. The data set, drawn from Basketball Reference, includes 346,785 player-games.

The most important and contentious feature of a CBA negotiation is financial: the share of league revenue to be paid in salaries, the salary cap which all teams will adhere to, and all accompanying compensation rules. This includes the value, length, and payment terms of player contracts. There are three contract types of interest which are available to players in free-agency (able to negotiate compensation with multiple teams). First is the minimum contract, which ensures that no player will receive less than \$1.1 million (in 2023-2024) per year. Veteran players earn larger minimum salaries, subsidized by the league: a player with 10+ years in the NBA by 2024 could not be paid less than \$3.2 million per season. The minimum contract is also the only contract freely available⁶ to teams whose payroll goes over the salary cap (~\$136m in 2023-2024) and is generally reserved for players who are at or near the replacement level or quality floor of the

⁶There are various and limited exceptions to this rule, including the mid-level exception.

league. While many minimum players provide good value to their teams, they are generally nearer to retirement or unemployment than to superstardom.

A second contract type is the Maximum contract, which pays a percentage of the salary cap. The specific percentage – 25, 30, or 35% (the “supermax”) – depends on the player’s years in the NBA and awards won in seasons prior to signing a new contract. The Maximum contract is almost exclusively reserved for the best players in the league (or players that teams project to become superstars), but it substantially underpays them. In 2013-2014, for example, Blake Griffin⁷ was paid just over \$16 million, but generated on-court performance that would have been valued, on an uncapped free-agent market, at around \$34 million⁸.

Finally, players who are not offered a maximum contract but are too good to play on a veteran minimum negotiate wages in between the two extremes. The market for this “middle-tier” of free agents is, by no means unfettered or efficient. Teams are constrained from spending freely by the salary cap and their prior salary commitments. Rookie contracts are fixed by draft position and early-career players have their bargaining power reduced via “Restricted Free Agency” which gives their current team considerable leverage over the player and against other bidders. Players also value other compensation besides money. Some prefer to live in Los Angeles or to play with a superstar, beloved coach, close friend, or blood relative. Still, this is a substantial set of transactions in which the market values a player’s performance and the primary place in which we can observe the wage impact of small changes in a player’s performance profile.

2.3. The Benefits of Professional Basketball as a Research Setting

Professional basketball offers three advantages for the study of worker performance and organizational outcomes. First, all organizations and workers are subject to a set of common rules and share (ostensibly) the same definition of success (Aversa et al. 2025). While there is certainly between-organization variation in resources, planning horizons, risk tolerance, and skill/sophistication,

⁷Go Clippers!

⁸Following FiveThirtyEight, I estimate a player’s open-market wage as the product of their wins (e.g., Win Shares) and the prevailing market price per win (approximately \$2.8 million in 2013–14).

incentives broadly push all teams in the direction of winning games and championships within a reasonable time frame. Even cases of “tanking” or losing games in order to obtain roster-building advantages in subsequent seasons are motivated by the eventual desire to win.

The second and third benefits derive from the fact that professional sports have rich and standardized performance data (Downs et al. 2025). Whereas many organizations are forced to hire, promote, or fire workers based on only a handful of observed data points, professional basketball teams can make decisions in an environment where every pass, shot, and step is logged. This greatly reduces the concern that decision-makers simply cannot see the distribution of a worker’s performance. Were we to observe promotion decisions made in a less information-rich environment, we would first have to question whether or not the decision-maker was even aware of the full breadth of their employee’s performance. And, as discussed, we would have every reason to think that they were not.

Finally, this wealth of data also allows researchers to partially observe the effort a player is exerting. In an oppositional game, more valuable actions are more costly. In the modern NBA, this means that the highest-value locations on the floor require the most energy to reach. A player seeking to shoot from 40 feet away can do so at any time (with the blessing of the opposing team). Higher-value shots (at the rim or from the free-throw line) require that the player be willing to work harder and, in all likelihood, absorb physical punishment in exchange for their shot attempt. Over the course of a season, a player’s willingness to run themselves to exhaustion or be elbowed/shouldered somewhere soft will naturally ebb and flow, which is reflected in the type of shots they seek out.

3. Study 1: Does the NBA value reliability?

The NBA is often called a “superstar’s league” because team marketing, fan engagement, and competitive success tend to follow the best players. As such, teams are sophisticated in measuring and compensating top talent. A player whose expected performance justifies a maximum contract will likely have several suitors and the oddities of the CBA mean that the intensive margin of compensation with respect to performance is nonexistent. The same is true, albeit for opposite

reasons, for players who are not good enough to earn above the minimum contract. If teams value the distribution of a player's performance, it should show up in the middle tier of contracts alone.

3.1. Methods and Data

Across the data set, I observe 2,347 free-agent contract signing events and classify them according to the CBA's wage scales. Of these, 201 are maximum contracts and 570 are minimum contracts (with both identified by the CBA's published wage scales). The remaining 1,576 are middle tier, and are identified by the percentage change in annual value. Contractual raises cannot exceed 8%, so any large increases must be achieved in free agency. I also include salary decreases of the same magnitude.⁹ This approach excludes players who re-signed at essentially unchanged salaries. For clarity, these three tiers are subsequently subscripted as c_{cap} , c_{floor} and c_{mid} .

At each contract signing event, the primary dependent variable is the player's log-transformed first-year salary on the newly signed contract. The independent variable is the player's volatility (noise-adjusted Game Score SD) in a window of 1, 2, or 3 prior seasons. For each tier, I regress compensation on a measure of the player's prior-period performance volatility, controlling for prior-period mean performance and a battery of career-stage and role covariates (some shot profiles are inherently more volatile than others). The regression includes season fixed effects. Standard errors are clustered by player. Formally, for player i 's contract in time t :

$$y_{i,t} = \beta_{\sigma} Volatility_{i,t-1} + \beta_{\mu} Mean_{i,t-1} + Controls_{i,t} + \delta_t + \epsilon_{it} \quad (3)$$

⁹I classify a signing as newly negotiated when its first-year salary differs from the player's prior salary by more than 9% (the 8% cap plus a small buffer for proration and reporting error).

Table 1: Variables and sample means by contract tier (signing events, 2011–2023)

Symbol	Definition	Mean by contract tier			
		Min	Mid	Max	All
y_{it}	First-year salary (\$M) of the newly signed contract; enters regressions in logs	1.30	7.08	29.04	7.56
$\text{Volatility}_{i,t-1}$	Noise-adjusted SD of per-36 Game Score over the evaluation window	4.73	5.08	5.69	5.10
$\text{Mean}_{i,t-1}$	Mean per-36 Game Score over the same window	9.74	11.41	15.81	11.67
age_{it}	Player age in signing year t	26.1	26.7	28.5	26.7
age_{it}^2	Squared age	—	—	—	—
$\text{tmin}_{i,t-1}$	Total minutes played, season $t - 1$; enters in logs	861	1,624	2,244	1,494
$\text{FGA}/\text{min}_{i,t-1}$	Field-goal attempts per minute, season $t - 1$	0.303	0.336	0.443	0.337
$\text{3PA}/\text{min}_{i,t-1}$	Three-point attempts per minute, season $t - 1$	0.101	0.105	0.127	0.106
$\text{FTA}/\text{min}_{i,t-1}$	Free-throw attempts per minute, season $t - 1$	0.073	0.086	0.133	0.087
δ_t	Signing-year fixed effects	—	—	—	—
N (signings)		570	1,576	201	2,347

The coefficients of interest are β_σ and β_μ . Both Table 2 and Table 3 report them in raw per-36 units: the change in log-wages associated with a one-point increase in per-36 mean Game Score or in the noise-adjusted per-36 standard deviation. Absent some bizarre circumstance, β_μ should always be positive, indicating that one additional point of Game Score is worth a higher salary. In raw units, the ratio $m_{mkt} = -\beta_\sigma/\beta_\mu$ represents the amount of mean Game Score a

team is implied to sacrifice, at a given wage point, in exchange for a one-point reduction in the noise-adjusted standard deviation of Game Score.

This analysis primarily examines the intensive margin, but volatility can also plausibly impact whether or not a player earns a contract of a given type or remains in the league at all. Using a multinomial logit model, I estimate the effect of increases to both mean performance and performance variance on the log-odds of a player's next free-agent contract falling into the middle or maximum tiers, rather than a minimum contract:

$$\ln \left[\frac{\Pr(\text{tier}_{i,t} = c)}{\Pr(\text{tier}_{i,t} = c_{\text{floor}})} \right] = \gamma_{\sigma,c} \text{Volatility}_{i,t-1} + \gamma_{\mu,c} \text{Mean}_{i,t-1} + \text{Controls}_{i,t} + \delta_{t,c}, \quad c \in \{c_{\text{mid}}, c_{\text{cap}}\} \quad (4)$$

I then use a simple linear probability model to test whether an at-risk player is re-employed. The sample keeps player-seasons of at least twenty games that are followed either by a new contract or by disappearance from the league; seasons whose continuation is contractually guaranteed are excluded, since those players mechanically cannot exit. This model uses the same set of controls and year fixed effects, and clusters standard errors by player.

3.2. Results

For negotiated, middle-tier contracts, more volatile players earn less. Regression 1 evaluates log-wages on the prior season's game-over-game volatility without controlling for mean performance and recovers a positive coefficient, suggesting that more volatile players earn more than reliable players. As in other disciplines, the best NBA players are also players who play the most, do the most, and so have the largest spread on performance. However, the sign on β_{σ} reverses once the mean of the player is included in Regression 2. Over a single season, the penalty is statistically indistinguishable from zero (Regressions 2 and 3). It emerges, and grows monotonically, as the evaluation window is expanded to two and then three seasons (Regressions 4 and 5).¹⁰

¹⁰Sample sizes fall as the evaluation window widens because more prior-season data is required for the analysis.

A one-point increase in per-36 volatility (equivalent to +0.94 SD) measured over the prior three seasons is associated with a statistically significant 8.9% decrease in wages on the player's next free-agent contract ($p = .001$). When measured over two prior seasons, the penalty is 4.0% ($p = .06$). This trend reflects the difficulty of estimating a worker's variance compared to their mean (β_μ is precisely estimated in every evaluation window). As virtually all free-agent contracts are signed after 2+ (and usually 4+) years in the league, the coefficient in Regression 5 most accurately represents organizational attitudes towards worker volatility.

Volatility is persistent, rather than the artifact of a single season. When measured two seasons before the signing, or one season after it, a player's performance standard deviation predicts the wage penalty (Before: -0.036 , SE 0.011 ; and after: -0.040 , SE 0.012 , per point), and, obviously, post-signing volatility cannot have caused a contract that was already signed.

Table 2: Reliability premium on log(wage). Coefficients in raw per-36 Game-Score units; standard errors in parentheses, clustered by player. Columns (6) and (7) re-estimate the column (5) specification within the fixed tiers. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	Negotiated Contracts					Fixed Contracts	
	(1)	(2)	(3)	(4)	(5)	(6) Min	(7) Max
β_σ (SD(GS))	+0.0318*** (0.0121)	-0.0049 (0.0125)	-0.0139 (0.0138)	-0.0412* (0.0223)	-0.0931*** (0.0290)	-0.0460 (0.0323)	-0.0141 (0.0307)
β_μ (Mean GS)	— —	+0.0594*** (0.0066)	+0.0519*** (0.0096)	+0.0646*** (0.0114)	+0.0910*** (0.0151)	+0.0240 (0.0163)	+0.0000 (0.0090)
age, age ²	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Log total minutes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Shot-attempt controls	No	No	Yes	Yes	Yes	Yes	Yes
Vol window (seasons)	1	1	1	2	3	3	3
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	1,176	1,176	1,176	925	677	77	123

The coefficients from Regression 5 are then used to generate estimates for the iso-value

slope m_{mkt} , or the inherent tradeoff between mean performance and volatility. All three evaluation windows have a positive value, suggesting that teams are willing to give up mean performance to get consistency at a given price point. However, the 95% confidence interval only excludes zero ([0.42, 1.63]) when teams have three seasons of prior performance to evaluate. Holding role, age, and shot profile constant, a value of 1.02 suggests that the market pays the same wage premium for a one-point increase in average performance as it does for a one-point reduction in volatility.

Table 3: Iso-value slope $m_{mkt} = -\beta_\sigma/\beta_\mu$, mid-tier signings. Raw Game-Score units; delta-method standard errors in parentheses; controls as in Table 2. * $p < 0.10$, *** $p < 0.01$.

Evaluation Window	β_μ	β_σ	m_{mkt}	95% CI	N
1 season	+0.0519*** (0.0096)	-0.0139 (0.0138)	+0.27 (0.26)	[-0.24, +0.78]	1,176
2 seasons	+0.0646*** (0.0114)	-0.0412* (0.0223)	+0.64 (0.34)	[-0.02, +1.30]	925
3 seasons	+0.0910*** (0.0151)	-0.0931*** (0.0290)	+1.02 (0.31)	[+0.42, +1.63]	677

Regressions 6 and 7 from Table 2 return the expected null results on the intensive margin. The minimum and maximum contracts offer pre-negotiated compensation, so a player who is already a max-worthy player does not earn more money by playing incrementally better nor is he punished for being more volatile. Table 4 reports the results of the multinomial logit and linear probability models which assess the relationship between extensive margin and performance volatility. Column 1 reports the difference in odds of signing a mid-tier contract versus a minimum tier contract. Column 2 compares maximum and mid-tier contracts. Column 3 reports the likelihood that a free-agent is re-employed in the subsequent season.

Unsurprisingly, the results for mean performance are positive and significant. A one-point increase in mean per-36 Game Score raises the odds of signing a middle-tier rather than a minimum contract by 21% and the odds of a maximum rather than a middle-tier contract by 30% (both

$p < .001$). However, the same increase in volatility has no significant effect.

Volatility likewise does not push players out of the league. Among at-risk player-seasons, re-employment rises by 1.9 percentage points per point of mean performance (SE 0.5) and by 1.6 percentage points per point of volatility (SE 0.8), suggesting that teams may be more likely to give one last chance to a volatile talent who briefly shows some flash of high-level performance.

Table 4: Volatility and the extensive margin. Columns (1)–(2): odds ratios per one-point increase from Equation (4), p -values in parentheses (MLE standard errors; column (2) is the maximum-versus-middle contrast, delta method). Column (3): linear probability of re-employment for free-agents, player-clustered standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	(1) Mid vs. Min	(2) Max vs. Mid	(3) Re-employment
Mean (per-36 GS)	1.21*** ($< .001$)	1.30*** ($< .001$)	+0.019*** (0.005)
Volatility (SD per-36 GS)	1.00 (.976)	0.87 (.158)	+0.016** (0.008)
Log total minutes	9.37*** ($< .001$)	8.16*** ($< .001$)	+0.381*** (0.036)
Career stage (age, age ²)	Yes	Yes	Yes
Shot-attempt controls	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
N	1,536	1,536	2,101

These results have two implications. First, they suggest that the results reported on the intensive margin are not the result of selection into the middle tier of contracts on the basis of a player’s performance spread. Second, they suggest that teams may evaluate performance sequentially. Initially, they assess how “good” a player is based on their mean performance. For players who are not worth a maximum contract but desirable enough to justify a mid-tier contract, teams then price volatility in conjunction with mean performance.

3.3. Discussion

The question of whether or not NBA teams value reliability has already been answered twice with conflicting results. Barnes and Morgeson (2007) suggest that there is a negative association between performance variability and wages. Awtrey et al. (2021) find no relationship. Neither analysis separates persistent, player-specific volatility from variation generated by playing time nor allows for the compensation nuances in the NBA CBA. Study 1 adjusts for these measurement and institutional constraints and shows a positive, significant relationship between a player's game-over-game reliability and their negotiated compensation. The association also emerges more strongly as teams accumulate multiple seasons of performance information, consistent with the greater difficulty of estimating volatility than mean performance.

This gives reason to be skeptical of any behavioral explanation for volatility's neglect in research and practice. When given information about worker volatility and given the option to use it, teams seem willing to price the added dimension of worker performance. The null results on the extensive margin still reinforce the intuition that teams first price a player's mean quality, then negotiate, where they can, for reliability.

It is important to note that the observed relationship is not causal. Thus far, I have made no effort to determine the compensation that a player would earn with their second-best offer or in some counterfactual in which they were more reliable but otherwise unchanged. Further, teams have more information about players than we do as fans or researchers. Perhaps there are personality or environmental traits which make volatile players unattractive for reasons other than their on-court performance. Teams with detailed medical information and more insight into a player's precise tendencies may simply value dimensions of health, availability, or fit which are unavailable in the game/performance data.

While I remain unable to rule out the presence of privately known and unobserved (in this data set) heterogeneity, many of these confounders could plausibly bias the observed results towards zero. Information, in professional basketball, is a source of competitive advantage and so does not flow freely. This first study suggests that teams will pay a premium for reliability, but teams

with information about some unattractive player characteristic will undervalue that player, relative to uninformed bidders. In all likelihood, the informed team will not sign the free-agent, leading players with hidden flaws to sort into uninformed teams. And it only takes one stupid, desperate, or risk-loving team to create a market for the headcase that 29 other franchises would prefer not to employ.

4. Study 2: Does reliability win basketball games?

Study 1 suggests that teams value reliability in the free agents they sign, but it says nothing about the normative value of that reliability. Nor does the question “should teams pay for steadiness” have an obvious answer. On one hand, professional basketball teams are generally “smart” and so we might assume that the decision to pay for something is rational and well-informed. On the other hand, research on organizations is replete with examples of mispriced attributes, biased preferences, and an over-adherence to received wisdom. We need not look further than professional sports, where baseball teams systematically overvalued batting average for generations (Hakes & Sauer 2006) and American football teams long resisted the statistically superior fourth down and two-point conversions (Hartley 2017; Romer 2006).

Even if we were convinced that streakiness is bad and consistency is good, it is possible that organizations have already evolved the tools and heuristics to manage it. In professional basketball, this might be the coach’s scheme or playbook. In the broader set of organizations, this would be handled with routines and best practices. Of course, this may not even be necessary. After all, even compared to other professional sports, professional basketball allows for a large number of draws. A game has an average of 100 possessions (compared to 35-40 plate appearances per team in Major League Baseball) and a season has 82 games (compared to 17 in the National Football League and 38 in the English Premier League). And, unlike other sports, every player can contribute to every possession, so the law of large numbers has ample opportunity to insulate team outcomes from individual player volatility.

4.1. Methods and Data

Whether or not they are correct to price reliability, teams are undeniably strategic actors. They select their rosters and allocate playing time with clear objectives. In order to address the endogeneity of these decisions, I exploit the quasi-random timing of injuries to key players. As a treatment, this has three useful features. First, while teams know that some players are more injury-prone than others, the precise timing of those injuries is difficult to predict. Further, teams use information about player injury risks strategically: fatigued players are given extra rest, prophylactic medical care, and tailored, state-of-the-art physical therapy, dietary advice, and equipment. Any injuries that result in spite of these measures are, I argue, an unpredictable residual. Finally, teams are heavily constrained in their mid-season roster choices, so injuries almost always entail the reallocation of minutes to players who are already on the roster (and whose prior performance we can measure). In effect, each injury involves two movements: the unexpected subtraction of a player and their replacement. Both pieces are known and measured, so I compare the team's quality and volatility before and after the injury against team outcomes. The estimand is the effect of the subtraction event as teams experience it, including whatever strategic response the team can muster. Results are substantially unchanged when directly controlling for the quality and volatility of the remaining roster.

I identify treatment events as player absences for which three conditions are met. First, the absent player must have averaged at least 15 minutes per game in the current season and played in each of his team's previous five games. This ensures that the injured players are active contributors (rather than experiments or fringe players) with sufficient performance history that I can examine pre-absence trends. Hereafter, I refer to these as "rotation players." Second, the player must play zero minutes for at least five games in a row.¹¹ This is a long enough spell that the team must genuinely reorganize its rotation in response. Finally, I keep only the first qualifying spell of

¹¹I do not use formal injury designations for two reasons. First, the data available is incomplete, with some verified injuries being reported as "Did Not Play - Coach's Decision" and others not reported at all. Second, there is nothing unique about injury as a reason for a player's absence compared to paternity leave, bereavement, or suspension for cause. For my purposes, these all result in the unexpected subtraction of a player from the team's roster for an indeterminate amount of time.

each team-season so that one injury event (and the corresponding roster/rotation changes) do not contaminate subsequent events. We might imagine that a team responds to an injury to its second-best player by asking the best player to do more. This added workload may cause a subsequent injury and make it impossible to disentangle the effects of the two injuries from each and from the impact of the shift in workload. By using only the first qualifying absence of the season, I avoid this altogether.

For each event, the focal player’s mean and volatility are measured over the prior season, using noise-adjusted, mean-centered Game Score and BPM. Game Score remains the primary trait of interest, but BPM is useful as it is a more robust measure of player mean quality in the long run. Players without a prior season (rookies, players who missed an entire season due to injury, players who were under contract but not good enough to play) are excluded. Across the 2011-2023 seasons, there are 268 qualifying events. Absent players average 27 minutes per game before their absence with a prior-season per-36 mean Game Score of 9.0 and volatility of 5.1. The median spell lasts 11 games.

I estimate the treatment effect of losing a player using a two-stage, staggered treatment difference-in-differences estimator (Butts & Gardner 2022; Gardner 2022). A conventional two-way fixed-effect design may produce biased estimates when treatment is staggered because some already-treated observations are used as controls for later cohorts. This estimator avoids this by ensuring that treated games are excluded from the counterfactual. I use Gardner (2022) instead of Callaway and Sant’Anna (2021) because the parameters of interest are coefficients on continuous moderators of the treatment effect. In the first stage (identical to Borusyak et al. 2024), I estimate team-season and game number fixed effects by regressing the result of team j ’s game g as follows:

$$y_{j,g} = \text{TeamSeasonFE}_j + \text{GameNumberFE}_g + \epsilon_{j,g} \quad (5)$$

on $\{(j, g) : \text{UntreatedGame}_{j,g} = 1\}$

This first stage uses games before and after the treatment period. I drop the last game before the absence because injuries which occur mid-game produce distorted statistics. I then estimate the

gap between predicted and actual team results for all games:

$$\hat{u}_{j,g} = y_{j,g} - \widehat{\text{TeamSeasonFE}}_j - \widehat{\text{GameNumberFE}}_g \quad \text{for all } (j, g). \quad (6)$$

In the second stage, I regress this gap against post-period indicators for player mean and standard deviation:

$$\hat{u}_{j,g} = \tau \text{Post}_{j,g} + \beta_\mu \text{Post}_{j,g} \text{Mean}_i + \beta_\sigma \text{Post}_{j,g} \text{SD}_i + \epsilon_{j,g}. \quad (7)$$

I use three outcome variables. One is a simple result indicator (1 if win, 0 if loss). A second is the point differential, where positive numbers indicate the margin of victory and negative numbers the margin of defeat. The third, sometimes called Pythagorean wins or excess wins, adjusts the binary result indicator based on the prior-season strength of the two teams playing and the location of the game (home/away).

Following Gardner (2022), which treats the two stages as a single estimation problem, I compute standard errors that add the uncertainty inherited from the first stage to the second-stage variance. I cluster at the team-season level because first-stage estimation error is shared by every game within a team-season.

The coefficients of interest are τ , representing the average treatment effect on the treated (ATT), and β_μ and β_σ , which measure how the ATT varies with changes to the absent player's mean and volatility.

4.1.1. Treatment Timing

Figure 3 illustrates the construction of the pre-period, treatment period, and counterfactual samples. The first row illustrates a below-median length absence. The second shows a player whose absence is protracted. The game in which the player is injured (region 2 in Figure 3) is not included in any samples because the player's statistics, in that game, are likely to be distorted by the mid-game injury.

I define the treatment-period as games 1-11 after the injury (region 3). This fixed window

is designed to include the entire time period in which a team is adjusting to the absence of a key contributor. The median injury length is 11 games. I prefer a fixed post-treatment evaluation window length because a player's return from injury is partially a team choice. Some players are given long periods of time to recover to full health. Others are rushed back into action for a key game. Were the study to look at results only after the injured player returns, it would be capturing the effect of the team's timing decision.

The design estimates the team and time fixed effects from untreated periods of time (regions 1,4, and 5), which is defined to include games before the player is injured and after the treatment window. Though this approach eases concerns about the endogeneity of a player's return, it is biased in multiple ways. For injuries that are shorter than the 11 game evaluation window, the "treatment" includes games where the injured player has already returned (region 3a in Figure 3). Games the player misses beyond the 11 game window (region 4) enter the counterfactual pool, contaminating a baseline that is meant to represent the team with him available.

Because the treated event is the first qualifying absence of each team-season, no other qualifying absence is in progress when the window opens. However, when another rotation player gets injured before the first returns (roughly half of the events overlap with another absence), I treat it as part of the effect being measured because the subsequent injuries could plausibly have been caused by overexertion in response to the first absence. Further analysis and discussion are provided in the appendix. Fortunately, the biases introduced by both short and long injury spells work against the reported ATT, making this a more conservative estimate of treatment effects than an alternative which uses the exact games in which the player is absent.¹² Games in Region 3a (player is available but assumed to be absent) are going to reduce the observed impact of their absence. Games in Region 4 (player is absent but assumed to be available) will lower the baseline and so reduce the observed difference between counterfactual and treatment periods. The same logic also explains why I do not control for other injuries to key rotation players outside the treatment window.

¹²Appendix includes estimates from alternative specifications.




Figure 3: Illustration of treated and untreated periods for short and long injuries.

4.2. Results

By using only injuries to rotation players as the treatment, this study virtually guarantees that all injuries hurt the team. A team plays its best players more than its worst players. When a player who receives significant playing time is absent, the team can only allocate those minutes to players that are worse (else they would have been playing in lieu of the lost player) or to players who are better but playing at capacity (else the team would already have allocated more minutes to them). Table 5 reports the expected negative and significant coefficient on the post-period indicator for all three outcome variables. Any injury costs the team, on average, about a point per game and just over two percentage points of win probability.

We would similarly expect a negative coefficient on β_μ . When a better player is injured, the team is hurt worse. Relative to the average injured rotation player, each additional point of prior-season Game Score costs the team about a sixth of a point of result margin and just over half of a percentage point of win probability. The positive sign on β_σ for all specifications suggests that an increase in the standard deviation of the injured player (holding mean performance constant) actually blunts the impact of their loss.

Table 5: Two-stage difference-in-differences estimates. Each row is a separate regression with different outcome variables. Two-step standard errors, clustered by team-season, in parentheses. $n \approx 21,500$ team-games; 268 treated cohorts. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	<i>Game Score</i>			<i>Box Plus Minus</i>		
	τ	β_μ	β_σ	τ	β_μ	β_σ
Point differential	-0.974*** (0.337)	-0.169* (0.090)	+0.612*** (0.189)	-0.974*** (0.334)	-0.319** (0.136)	+0.380** (0.154)
Win rate	-0.022* (0.012)	-0.0052* (0.0031)	+0.0222*** (0.0072)	-0.022* (0.012)	-0.0088* (0.0046)	+0.0124** (0.0055)
Excess wins vs. Pythagorean	-0.023** (0.012)	-0.0053* (0.0030)	+0.0208*** (0.0070)	-0.023** (0.012)	-0.0093** (0.0045)	+0.0123** (0.0053)

To illustrate the magnitude of these results, consider four hypothetical players (Table 6). Player A is the average injured rotation player in the sample, with a per-36 mean of 9 and volatility of 5. Because the moderators are mean-centered, the cost of losing him is about one point of margin per game, or a quarter of a win over the median eleven-game absence. Player B is a near-star whose absence clearly hurts the team. Player C is the steady performer whose absence is greatly missed. Player D illustrates that sufficient volatility can completely offset the win-loss contributions of a league-average rotation player.

Table 6: Predicted cost of absence for four hypothetical players (Game Score coefficients).

	Mean (GS)	Volatility (GS SD)	Point differential per game	Losses (per-11 games)
A: average player	9	5	-1.0	+0.24
B: near-star	13	5	-1.7	+0.47
C: one point steadier	9	4	-1.6	+0.49
D: one point streakier	9	6	-0.4	0.00

4.2.1. Causality and Robustness

Figure 4 presents an event plot in which each bucket holds four games. In the four games before the injury and the first four afterwards, team point differential is statistically indistinguishable from zero (joint Wald $p = 0.49$ and $p = 0.16$ respectively). One possible explanation is that teams take immediate tactical and strategic steps to close the gap. However, the increased load on the remaining rotation (plus the opportunity for opposing teams to respond strategically) means that the relief is only temporary. In the subsequent 8 games, team performance falls significantly before rebounding in the final two buckets. The figure also splits the events into two cohorts above and below median volatility, suggesting that the ATT is largely driven by injuries to more reliable players.

The estimates in Table 5 support a causal story only if the games in the post-injury evaluation window would otherwise have resembled the team’s season-long performance. More specifically, causality is threatened if there is something about the treatment window that impacts team performance, is correlated with injuries, but does not result from player absences. This is distinct from other types of event studies which use the pre-treatment period to estimate the counterfactual and so require support for the parallel trends assumption.

I examine several plausible candidates. I first examine whether overexertion could be responsible. If injuries follow stretches of unsustainably high-effort play, the subsequent results would be poor with or without absences. However, recent performance does not predict when an absence begins ($p = 0.3$ across 4,886 team-games played before the first qualifying absence of the season) and the measured effects are unrelated to the team’s pre-injury performance (β_σ moves from 0.61 to 0.59 when an explicit pre-event control is added). Likewise, there is no significant difference between treated and untreated games on location (home/away), opponent strength, or rest days between games. Finally, teams late in the season may change priorities (sometimes called tanking) from winning games to developing younger players or trading rotation players for other assets. The study design nearly rules this effect out by construction since only 5 of the 268 events begin at or after game 55.

Placebo tests using games before the injury or after the injured player returns have negligible coefficients for margin of victory/defeat (pre-injury $\beta_\sigma = +0.01$, post-return $\beta_\sigma = +0.03$). Likewise, when the injured player's characteristics are replaced with those of a different, randomly selected player, estimates exceeding those in Table 5 are only reproduced in 20 out of 2,000 re-runs using BPM. For Game Score, it is 4 out of 2,000. Taken together, these robustness checks suggest that whatever drives β_σ operates only while the player is not playing.




Figure 4: Team point differential around a rotation player's absence, in four-game buckets (the injury game, game 0, is omitted). 95% intervals from two-step standard errors, clustered by team-season.

4.2.2. On Steadiness and Superstars

The results presented in Table 5 are largely driven by players who excel on one or more dimensions. Figure 5 reports the binned ATT by quintile of prior season mean or volatility. The bins are balanced on playing time. Where previous plots use Game Score, I now shift to using Box Plus Minus (BPM) because it is a more sophisticated measure when making player-to-player comparisons. BPM's game values are estimated at the season level and are not interpreted here. However, the game-to-game dispersion is at least as reliable a measure of spread as Game Score's. Once more, negative values indicate the estimated harmful impact to team performance from losing a player with mean or

standard deviation in a given quintile. The results for mean are intuitive. Losing below average players does not hurt the team nearly as much as losing a very good player (BPM above 4.0 is considered “All-Star Quality” per Myers (2020)).

Likewise, the loss of one of the most reliable players hurts team point differential. What is somewhat surprising is the fact that the magnitudes of the two effects are somewhat similar. Losing a star player costs the team 2.64 points per game. Losing a very reliable player costs 2.14 points.




Figure 5: Post-injury ATT on point differential by minutes-balanced quintile of the absent player’s mean or SD of prior-season BPM. Error bars are $\pm 1.96 \cdot$ bootstrap SE.

Interestingly, these appear to be almost completely different populations of players. Figure 6 plots the observed player absences and highlights the players in each of the top quintiles from Figure 5. The overlap is tiny. There are only four players (out of 268) with observed absences who are both all-star quality and highly reliable. Of course, this is not to say that players cannot be both. Rather, the rare player who is a highly reliable superstar is also unlikely to get injured early in the season. Still, evidence suggests that Reliable (triangles) and High-Average (squares) are two different populations with roughly similar marginal replacement costs. Crucially, the former delivers 80% of the impact while being paid, on average, 40% as much as the latter group (Table 7).




Figure 6: Absent or injured players by prior-season mean BPM and prior-season SD of BPM. Triangles represent players in the top quintile of Reliability. Squares are the top quintile of mean performance. Stars are in both categories.

Table 7: The price of losing a player: minutes-weighted ATT on point differential by quintile of the absent player's prior-season BPM with salary and share of salary cap. Bootstrap standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	ATT (pts/game)	Salary	% of cap
<i>Panel A: quintiles of mean BPM</i>			
Q1 (lowest mean)	-0.03 (0.67)	\$4.5M	5.8
Q2	+0.25 (0.96)	\$7.9M	9.1
Q3	-1.70*** (0.64)	\$10.0M	11.1
Q4	-1.74** (0.85)	\$12.1M	13.4
Q5 (near-stars)	-2.64*** (0.89)	\$18.7M	20.0
<i>Panel B: quintiles of SD of BPM</i>			
Q1 (most reliable)	-2.14** (0.84)	\$7.5M	9.0
Q2	-0.62 (0.63)	\$10.9M	12.8
Q3	-0.98 (0.83)	\$9.1M	11.3
Q4	-0.84 (0.94)	\$11.5M	12.1
Q5 (most volatile)	-1.29* (0.78)	\$12.4M	12.6

4.3. Discussion

Study 2 suggests that volatility is genuinely detrimental to team performance and outcomes. Even after the best efforts of coaches and teammates, who have more information about players' roles and tendencies and the ability to adjust schemes and rotations strategically, there is still a persistent and negative effect when a reliable cornerstone of the team is lost to injury. This validates the finding, from Study 1, that teams price reliability as a distinct player performance characteristic. The coefficients in Table 3 are consistent with those in Table 5. Teams trade roughly one point of mean performance for a one-point reduction in Game Score standard deviation, while the Study 2 coefficients value that same reduction at 3.6 points of mean production, though the ratio is imprecise (95% CI -0.5 to 7.8), so the price the market pays and the value of the production are statistically indistinguishable.

In the real world, superstars are rare. Not every company can have one, let alone an army of 10x engineers or A+ sales reps. And for companies with limited resources, top-tier human capital is often out of reach. These results suggest that those organizations may find that it is better to build a steady and consistent, if unspectacular, workforce, rather than compete with the deep pockets of the world's largest and most prestigious firms. This is not to say that superstars are not important (most of the aforementioned highly reliable players share a roster with one or more max-contract superstars), but for companies who are constrained relative to their competitors in the market for talent, reliability has a quality all its own.

Missing, from Study 2, is more detail on the exact mechanism that drives the results. After all, I have shown that players who consistently score fewer points (or do fewer things) can be, on average, better for the team than those that score in bunches with fallow periods in between, which is, perhaps, counterintuitive. As this is a paper about organizations, performance, and strategy and not a paper about professional basketball, the specific basketball mechanisms are largely outside the scope of Study 2. However, I can note three mechanical features that generalize beyond professional basketball.

First, success and failure are binary and separated by a knife's edge. The first point that

the opponent cannot match is extremely valuable (the team wins by 1) but the second point of win-margin is worthless (a win by two points is the same as a win by one point).¹³ In effect, the team is forbidden from keeping any excess performance past some threshold. This is a common feature of organizational payoff functions and suggests areas where reliability or safety should be valued: Political parties need only win by one vote. Entrants need only move fast enough to erect entry barriers behind them. And startups need only enough revenue to pay the bills in order to live to fight another day. In environments with the opposite incentives (limited or fixed downside, lots of upside), we would expect to see a premium on volatility (venture capital investing, computational/algorithm races).

Second is the shape of worker performance itself. Game-level output is right-skewed for roughly 90% of the rotation players in professional basketball. A player's best decile of games averages 1.9 standard deviations above his own mean, while his worst decile averages only 1.5 below.¹⁴ Volatile players, then, do not trade good games for bad games, one for one. They are trading many games below average for a handful that are well above average, with team payoffs again, limited by the nature of binary win-loss outcome.

Finally, as is true in many organizations, basketball is a team sport in which a player's contributions affect both the scoreboard and the contributions of their peers. A reliable teammate, in basketball, helps the offensive and defensive schemes function even when other players are out of position. They provide insurance that covers for mistakes or allows others to take risks without fear of outsized negative consequences. In settings where performance is entirely independent, we should see a weaker preference for reliability (as in professional baseball, which values power hitters highly, even if they strike out often). And where teamwork matters, it is worth paying a premium for consistency.

¹³Though professional basketball has flirted with rule changes which would modify this, including the Elam ending's target-score finish.

¹⁴Player-seasons with at least forty appearances and fifteen minutes per game, 2011–2023.

5. Conclusion

This paper suggests that organizations do price worker volatility consistent with its impact on team outcomes. Where the CBA allows prices to move, teams pay a premium for consistency. As evidence of a player's volatility accumulates, teams pay less for the same mean performance. When an injury removes a player, the cost of his absence rises with his steadiness (holding average performance fixed) and is consistent with the wage the market pays.

So why is worker volatility so conspicuously absent from strategic human capital management? Perhaps worker volatility doesn't matter to organizational outcomes, either because it is absorbed or buffered by existing structures and routines (Orton & Weick 1990; Weick 1976) or because the law of large numbers provides statistical insurance against extreme draws. Or, we might instead take the lack of attention paid to worker volatility as evidence of a managerial (and researcher) blind-spot. After all, individuals are known to struggle to separate signal from noise or draw the correct conclusion from stochastic signals (Aldrich & Ruef 2018; Benson et al. 2019; Denrell 2005; Lazear 2004; Liu & Tsay 2021; March & Sutton 1997). Finally, an organization-specific reading of Morgan's Canon¹⁵ suggests that we not attribute to either strategy or bias what can be better explained by measurement error. That is: the shape and tails of a worker's performance distribution are considerably more difficult, expensive, and time-consuming to measure than the mean.

These three explanations are neither collectively exhaustive nor mutually exclusive and the correct answer is almost certainly "a little bit of all of them." However, Study 2 rejects the strongest version of the first candidate. In a setting with ample strategic behavior and substantial scale (100 possessions per game, 82 games per season, each with 10 workers participating), the normative effect remains. Losing someone consistent produces an outsized and negative effect on team performance. This is not to say that buffering or formal routines and structures are irrelevant, but it does suggest that they are incomplete and costly enough to present teams with a dilemma.

¹⁵"In no case is an animal activity to be interpreted in terms of higher psychological processes if it can be fairly interpreted in terms of processes which stand lower in the scale of psychological evolution and development."

Likewise, the law of large numbers cannot fully average away outlier performances in a competitive setting with fine margins. Study 1, meanwhile, argues against the idea that decision-makers have a blind-spot. Teams price volatility where the market allows, though Table 7 suggests that the pricing may be incomplete.

What remains is measurement. While there is substantial evidence that volatility is a semi-stable worker characteristic, there is widespread disagreement about the shape of the distribution from which performance is drawn (Beck et al. 2014; O’Boyle Jr. & Aguinis 2012). Without knowing whether a worker’s performance is normal, log-normal, or heavy-tailed, how much data must an employer collect before confidently calling that worker reliable or volatile? And for researchers, attempting to simultaneously estimate both moments of each worker’s performance distribution is likely to strain most field and lab studies. Even in this setting, more than half of observed variance is mechanical noise and the wage penalty is only visible in the data over multi-season evaluation windows. But, where measurement is available, rational and strategic behavior follows.

In strategic factor markets, firms can capture value when they have more accurate expectations than competitors about the future value of a resource or the evolution of the world state (Barney 1986), which makes the generation, gathering, and interpretation of information into important capabilities for sustained competitive success (Makadok & Barney 2001; Ryall & Sorenson 2025). Though people are the eponymous human resource, strategic human capital would benefit from developing on these well-understood ideas. After all, worker reliability is important to organizational success, difficult to know ex-ante, costly to estimate, and not fully reflected in labor-market prices. Organizations that measure it more accurately may have a substantial and sustained advantage over their competition.

The immediate problem is one of measurement. Research on human resources should either have a credible reason for reducing workers to the mean or take steps to ensure that interventions which raise the average do not materially and adversely impact the distribution of worker performance. Measurement, in turn, gives way to management, which is important because workers do not produce averages. They deliver noisy, messy sequences of good and bad days, and the

organization gets to experience them all.

References

- Aldrich, H. E., & Ruef, M. (2018). Unicorns, Gazelles, and Other Distractions on the Way to Understanding Real Entrepreneurship in the United States. *Academy of Management Perspectives*, 32(4), 458–472.
- Aral, S., Brynjolfsson, E., & Wu, L. (2012). Three-Way Complementarities: Performance Pay, Human Resource Analytics, and Information Technology. *Management Science*, 58(5), 913–931.
- Aversa, P., P. Moliterno, T., Rockmann, K. W., Bothner, M. S., Sharapov, D., & Eckardt, R. (2025). EXPLORING NEW PLAYING FIELDS: SPORTS SETTINGS AND MANAGEMENT THEORIZING. *Academy of Management Discoveries*, 11(2), 145–151.
- Awtrey, E., Thornley, N., Dannals, J. E., Barnes, C. M., & Uhlmann, E. L. (2021). Distribution neglect in performance evaluations. *Organizational Behavior and Human Decision Processes*, 165, 213–227.
- Barnes, C. M., & Morgeson, F. P. (2007). Typical Performance, Maximal Performance, and Performance Variability: Expanding Our Understanding of How Organizations Value Performance. *Human Performance*, 20(3), 259–274.
- Barney, J. B. (1986). Strategic Factor Markets: Expectations, Luck, and Business Strategy. *Management Science*, 32(10), 1231–1241.
- Battiston, D., Espinosa, M., & Liu, S. (2025). Talent Poaching and Job Rotation. *Management Science*, 71(4), 2975–2992.
- Bavafa, H., & Jónasson, J. O. (2021). The Variance Learning Curve. *Management Science*, 67(5), 3104–3116.
- Bavafa, H., & Jónasson, J. O. (2024). The Distributional Impact of Fatigue on Performance. *Management Science*, 70(5), 3319–3337.
- Beck, J. W., Beatty, A. S., & Sackett, P. R. (2014). On the Distribution of Job Performance: The Role of Measurement Characteristics in Observed Departures from Normality. *Personnel Psychology*, 67(3), 531–566.
- Becker, B., & Gerhart, B. (1996). The Impact of Human Resource Management on Organizational Performance: Progress and Prospects. *The Academy of Management Journal*, 39(4), 779–801.
- Benson, A., Li, D., & Shue, K. (2019). Promotions and the Peter Principle*. *The Quarterly Journal of Economics*, 134(4), 2085–2134.
- Bidwell, M. (2011). Paying More to Get Less: The Effects of External Hiring versus Internal Mobility. *Administrative Science Quarterly*, 56(3), 369–407.
- Borusyak, K., Jaravel, X., & Spiess, J. (2024). Revisiting Event-Study Designs: Robust and Efficient Estimation. *The Review of Economic Studies*, 91(6), 3253–3285.
- Butts, K., & Gardner, J. (2022). Did2s: Two-Stage Difference-in-Differences. *The R Journal*, 14(3), 162–173.
- Callaway, B., & Sant’Anna, P. H. C. (2021). Difference-in-Differences with multiple time periods. *Journal of Econometrics*, 225(2), 200–230.

- Chadwick, C., & Dabu, A. (2009). Human Resources, Human Resource Management, and the Competitive Advantage of Firms: Toward a More Comprehensive Model of Causal Linkages. *Organization Science*, 20(1), 253–272.
- Coff, R. W. (1997). Human Assets and Management Dilemmas: Coping with Hazards on the Road to Resource-Based Theory. *The Academy of Management Review*, 22(2), 374–402.
- Coff, R. W. (1999). When Competitive Advantage Doesn't Lead to Performance: The Resource-Based View and Stakeholder Bargaining Power. *Organization Science*, 10(2), 119–133.
- Dalal, R. S., Alaybek, B., & Lievens, F. (2020). Within-Person Job Performance Variability Over Short Timeframes: Theory, Empirical Research, and Practice. *Annual Review of Organizational Psychology and Organizational Behavior*, (7), 421–449.
- De Figueiredo, R. J. P., Meyer-Doyle, P., & Rawley, E. (2013). Inherited agglomeration effects in hedge fund spawns. *Strategic Management Journal*, 34(7), 843–862.
- Denrell, J. (2005). Should We Be Impressed With High Performance? *Journal of Management Inquiry*, 14(3), 292–298.
- DeVaro, J. (2006). Strategic promotion tournaments and worker performance. *Strategic Management Journal*, 27(8), 721–740.
- Downs, J., MAHONEY, J. T., & SOMAYA, D. (2025). DID CHEATING HELP THE HOUSTON ASTROS WIN? ORGANIZATIONAL MISCONDUCT, ILLICIT COMPETITIVE INTELLIGENCE, AND ORGANIZATIONAL PERFORMANCE. *Academy of Management Discoveries*, 11(2), 228–254.
- Gardner, J. (2022). Two-stage differences in differences. *Papers*.
- Gilovich, T., Vallone, R., & Tversky, A. (1985). The hot hand in basketball: On the misperception of random sequences. *Cognitive Psychology*, 17(3), 295–314.
- Hakes, J. K., & Sauer, R. D. (2006). An Economic Evaluation of the Moneyball Hypothesis. *Journal of Economic Perspectives*, 20(3), 173–186.
- Hartley, J. (2017, February). Going for Two: Optimizing between Extra Points and Two Point Conversions in the NFL.
- Hollinger, J. (2003). *Pro basketball prospectus*. Brassey's.
- Lazear, E. P. (2004). The Peter Principle: A Theory of Decline. *Journal of Political Economy*, 112(S1), S141–S163.
- Liu, C., & Tsay, C.-J. (2021, October). The Variance of Variance.
- Lu, G., Du, R. Y., & Peng, X. (2022). The Impact of Schedule Consistency on Shift Worker Productivity: An Empirical Investigation. *Manufacturing & Service Operations Management*, 24(5), 2387–2796.
- Makadok, R., & Barney, J. B. (2001). Strategic Factor Market Intelligence: An Application of Information Economics to Strategy Formulation and Competitor Intelligence. *Management Science*, 47(12), 1621–1638.
- March, J. G., & Sutton, R. I. (1997). Crossroads—Organizational Performance as a Dependent Variable. *Organization Science*, 8(6), 698–706.
- Myers, D. (2020). About Box Plus/Minus (BPM).
- O'Boyle Jr., E., & Aguinis, H. (2012). The Best and the Rest: Revisiting the Norm of Normality of Individual Performance. *Personnel Psychology*, 65(1), 79–119.
- Orton, J. D., & Weick, K. E. (1990). Loosely Coupled Systems: A Reconceptualization. *The Academy of Management Review*, 15(2), 203–223.
- Perrow, C. (1984). *Normal Accidents: Living with High-Risk Technologies*. Basic Books.

- Podsakoff, N. P., Spoelma, T. M., Chawla, N., & Gabriel, A. S. (2019). What predicts within-person variance in applied psychology constructs? An empirical examination. *Journal of Applied Psychology, 104*(6), 727–754.
- Romer, D. (2006). Do Firms Maximize? Evidence from Professional Football. *Journal of Political Economy, 114*(2), 340–365.
- Ryall, M. D., & Sorenson, O. (2025). The Causal Inference Problem: When Can Managers Use Data to Inform Decisions and How Can Organization Design Help Them? *Academy of Management Review, 1*–25.
- Sackett, P. R., Zedeck, S., & Fogli, L. (1988). Relations between measures of typical and maximum job performance. *Journal of Applied Psychology, 73*(3), 482–486.
- Weick, K. E. (1976). Educational Organizations as Loosely Coupled Systems. *Administrative Science Quarterly, 21*(1), 1–19.