ECONOMIC RECORD

ECONOMIC RECORD, 2022

Superstar Productivity and Pay: Evidence from the Australian Football League

PATRICK J. FERGUSON and MATTHEW PINNUCK
Faculty of Business and Economics, The University of Melbourne, Carlton, VIC, Australia

We use game-level data from the Australian Football League (AFL) to examine superstar workers' productivity and pay. By exploiting teams' injury-induced line-up changes between games, we show that, compared with replacement-level players, superstars increase their teams' likelihood of winning away games by approximately 15 percentage points. While we then show that betting markets appear to appropriately price superstars' marginal productivity, we present back-of-the-envelope calculations that suggest that teams underpay superstar players by at least 30 per cent. We discuss how inaccurate performance evaluations, labour market regulations, long-term back-loaded contracts, clubs' attempts to reduce harmful intra-team pay disparities and on-field success as a form of payment-in-kind may explain our findings.

I Introduction

At the end of the 2013 Australian Football League (AFL) season, Hawthorn's Lance Franklin, at the time the league's most highly regarded forward, signed a 9-year, A$10 million contract to join the Sydney Swans. Upon its announcement, the deal was met with public outrage and media scrutiny. Beyond the standard moral objections to the level and distribution of pay in professional sports, critics primarily questioned whether the deal was efficient in an economic sense.[1] In the popular press, headlines declared that Franklin was 'overpaid, overrated, and over the hill' (Sydney Swans, 2013). Senior executives at the AFL described the deal as an 'extraordinary risk', one that, given the constraints imposed by the league's salary cap, would compromise the Swans' ability to remain competitive toward the end of Franklin's career (Herald Sun, 2013). The response to Franklin's contract was far from unusual. In professional sports more generally, team executives, administrators and regulators, player agents, pundits and casual fans have long speculated about the magnitude of the impact that superstars have on their teams' performance and the implications this has for not only player pay, but also team management and recruitment practices, behaviour in betting and spectator markets, and optimal trade and draft policies.

Building on Rosen (1981), the economics literature on superstars sheds some light on these questions.[2] A stream of this research provides

JEL classifications: J30, L83, M12, M52, Z22

Correspondence: Patrick J. Ferguson, Faculty of Business and Economics, The University of Melbourne, 198 Berkeley St., Carlton, VIC 3053, Australia. Email: pfer@unimelb.edu.au

[1] Critics challenged the size and duration of Franklin's contract. At the time, the popular opinion was that Franklin would likely be worth A$1 million per season in the early years of his contract, but unlikely to continue to be worth this toward the end of the deal. In Section 10IV.iv, we discuss the use of long-term contracts in the AFL, and how these types of deals may be efficient under certain conditions.

[2] Using a model that interacts human capital with production technology, Rosen (1981) describes the phenomenon of superstars, or the kinds of economic activity where there is concentration of output among a few individuals, marked skewness in the associated distributions of income and very large rewards at the top. Subsequent related research has explored how extreme value theory and power laws may explain the distributions of talent and pay in a range of different settings (Gabaix & Landier, 2008; Gabaix, 2016).

doi: 10.1111/1475-4932.12665

insightful descriptive evidence from a range of settings on the relationship between superstars' productivity and pay in an effort to understand whether large compensation packages reflect optimal incentive design or simply rent extraction on the part of entrenched or high-profile employees (Gabaix & Landier, 2008; Malmendier & Tate, 2009). Typically focused on examining major North American and European sports competitions, another stream of this literature investigates the impact that superstar athletes have on a variety of outcomes including stadium attendance and broadcast audiences (Hausman & Leonard, 1997; Berri et al., 2004; Jane, 2016; Jewell, 2017). Research in sports economics also discuss whether players are paid their marginal revenue products (Scully, 1974; Rosen & Sanderson, 2001) and investigates whether pay inequality within teams impacts team productivity (Depken, 2000; Breunig et al., 2014).[3]

While well-developed, empirical research in economics on superstars is primarily descriptive, and we are unaware of a study in this literature that employs an explicit identification strategy to isolate the causal effect of superstars on team performance. As further motivation for our study, we also note that despite the economic and cultural significance of Australian-rules football in this country (Blainey, 2010), prior research in sports economics has yet to examine superstar productivity and pay in the AFL.

In this paper, we address these gaps in the economics literature on superstars by exploiting plausibly exogenous variation in firms' workforce composition to identify workers' marginal productivity. To do so, we use injury-induced changes to AFL team line-ups between games to produce robust estimates of superstar players' marginal productivity. We then interpret these estimates alongside data on the level and distribution of player pay in the AFL. Specifically, we define superstars as players in the 95th percentile of the distribution of the league's most-widely used performance rating. To measure workers' marginal productivity, we compare the attributes and outcomes of games where due to injury superstars were withdrawn from their teams' line-up to the attributes and outcomes of games where superstars were not withdrawn due to injury. We show that as these injury shocks are plausibly exogenous, our research design allows us to cleanly identify superstars' marginal productivity. This provides an advantage over prior empirical studies on superstar productivity and pay, as these studies typically employ field settings where worker output is difficult to measure and likely to be confounded by firm outcomes and endogenous personnel decisions.

Using linear ordinary least squares (OLS) to implement this empirical strategy, we find that compared with replacement-level players, superstars in the AFL increase their teams' likelihood of winning away games by approximately 15 percentage points. Put differently, we show that compared with replacement-level workers, superstars improve their teams' margin in away games by approximately 15 points, or 2.5 goals.[4] We note that for home games these superstar effects are not statistically significant. We discuss potential explanations for this finding. First, we explore why our empirical strategy may be underpowered for statistical tests of superstar effects in home games. Second, we consider whether superstars effects are more pronounced under more challenging work conditions. Specifically, we discuss whether superstar players are especially effective at neutralising opponents' home-ground advantage and whether stars are subject to lower travel costs.

Next, we show that betting markets appear to appropriately price superstars' marginal productivity. Using changes in pre-game betting odds, we find that compared with replacement-level players, superstars in the AFL increase their teams' expected likelihood of winning home (away) games by approximately 11 (6) percentage points. We show that compared with replacement-level players, superstars increase their teams'

[3] A related body of research also studies 'pay' and college athletes' marginal revenue products (Kahn, 2007). This work is part of a larger literature in economics on monopsony in labour markets.

[4] These are estimates of the linear effect of a superstar on team performance. As we observe very few instances of a team losing more than one superstar to injury, we are unable to identify whether the effects of superstars are non-linear, for example, we cannot say whether the marginal effect of a second superstar is less than 15 percentage points. However, given both that superstars by definition are rare and that the AFL implements strict salary caps and free-agency restrictions that broaden the distribution of playing talent across the competition, teams in the AFL are only likely to have at most two superstars on their list. As such, understanding the effect of a second, third or fourth superstar on team performance is an exercise that has little practical meaning in our setting.

expected margin in home (away) games by approximately 13 (8) points.[5]

While we document that betting markets appear to appropriately price superstars' marginal productivity in the aggregate, we also present back-of-the-envelope calculations that suggest that AFL teams underpay superstar players by at least 30 per cent. We show that compared with replacement-level players superstars generate an additional 1.65 wins per season for their team. We then estimate that AFL teams spend on average approximately A$1,100,000 on player payments per win. Using these figures to calculate the value of superstars' marginal productivity, we find that superstars should be paid approximately A$2,100,000 per season. We compare this estimate with the distribution of player pay in the AFL and show that the 95th percentile of players are paid at least A$800,000 per season, while no player is paid more than A $1,500,000 per season. As we only estimate superstars' marginal on-field productivity and do not take into consideration superstars' impact on team revenue via increased attendance, membership, and merchandising sales, we argue that in practice superstars are likely to be more heavily underpaid than our back-of-the-envelope calculations suggest.[6]

We discuss several explanations for why superstars in the AFL may appear to be underpaid. First, we consider how inaccuracies in managers' performance evaluations may cause teams to underestimate the productivity of star players relative to the productivity of marginal players. Second, we discuss how managers may deliberately underpay superstars so as to mitigate the harmful effects of pay disparities among coworkers on team morale and performance. Third, we consider how labour market regulations may compress the pay distribution in the AFL and force teams to pay superstars less than the value of their marginal product. Fourth, noting the widespread use of backloaded contracts in the AFL that often see players underpaid in their prime and overpaid at the end of their careers, we discuss how superstars may not in fact be underpaid when we take into consideration players' total earnings over the duration of their contract. Finally, we discuss how, if players derive utility from monetary and non-monetary factors, superstars may accept on-field success, that is, winning premierships, as a form of 'payment in kind'.

As our findings hinge on the validity of our research design, we conduct a number of robustness checks that evaluate the plausibility of our empirical approach. To identify superstars' marginal productivity, injury-induced line-up changes between games must be plausibly exogenous. To support this assumption, we show that withdrawals of star players from the home and away teams' line-ups are unrelated to a range of game and team attributes. We also perform a pair of placebo tests. First, we conduct analyses whereby for each game in our sample we repeatedly assign injury-induced line-up changes to teams' star players at random and then show that the mean of these placebo treatment effects is non-distinguishable from zero. Second, we weaken our definition of superstar players and then show that we fail to observe superstar effects for players in the 90th–95th percentiles, that is, 'good-but-not-great' players. We motivate these tests using the logic that our main results are more likely to be spurious if we also observe superstar effects for these 'non-superstars'. This robustness check also provides insight into the question of 'who' should be considered a superstar in the context of the AFL.[7]

Finally, we also show that our superstar measure co-moves with four established measures of superstar performance: Best and Fairest votes, Brownlow Medal votes, All Australian team selection, and AFL Most Valuable Player votes. We argue that these descriptive findings support the overall validity of our approach for identifying superstar players in the AFL and strongly suggest that our main results are unlikely to be driven by idiosyncrasies of our specific superstar measure.

[5] For away games, our market-based estimates of superstar productivity are statistically weak at the two-tail level but reasonably strong at the one-tail level, assuming strong priors for a positive superstar effect.

[6] To make this claim, we need to assume that on- and off-field productivity are positively correlated (i.e., better players are more popular with fans) and that teams value both the former and latter outputs.

[7] While it would be insightful to estimate the marginal productivity of players at the 99th percentile, we lack the necessary number of observations to do so with any reasonable degree of precision, that is, we simply do not observe enough injuries among this very small group of elite players to implement our empirical strategy.

Overall, our study primarily contributes to the literature on the economics of superstars (Rosen, 1981). Malmendier and Tate (2009) and Groysberg et al. (2008) study superstar chief executive officers (CEOs) and security analysts, respectively. Azoulay et al. (2010) estimates that superstar life scientists generate large positive network effects that serve to heighten their collaborators' research output. In sports settings, economics research finds that superstar pay is highly convex and that superstar players are associated with large attendance externalities (Lucifora & Simmons, 2003; Humphreys & Johnson, 2020). Examining superstar pay in a local setting, Blackham and Chapman (2004) estimate that cricketer Donald Bradman had a large effect on Test crowd sizes and that this effect generated additional revenue for the Australian Cricket Board. We extend the economics literature on superstars by exploiting plausibly exogenous variation in firms' workforce composition to generate unbiased estimates of superstars' marginal productivity. We show in a field setting where employees face high-powered performance incentives and highly selective hiring practices that output is heavily concentrated among a subset of highly talented workers.

Our study also contributes to research on the economic behaviour of professional sports organisations. In particular, we extend a stream of this literature that examines the business of Australian-rules football, the country's most-popular and lucrative sport. In 2019, the AFL's consolidated revenue was approximately A$1.55 billion (AFL, 2019). For Australian broadcasters, AFL games are the most popular form of programming and a major source of advertising revenue (Dang et al., 2015). As the first study on the financial performance of clubs in the AFL, Pinnuck and Potter (2006) shows that teams' off-field financial success (i.e., match attendance and club memberships, etc) is positively related to both their short- and long-term on-field success.[8] Examining player payments in the AFL, Booth et al. (2012) finds that the player salary share of AFL expenditure and revenue data fell over the period from 2001 to 2009. We contribute to this literature by linking unbiased estimates of superstar players' marginal productivity to the distribution of player pay in the AFL. Ours is the first study to show that while superstars have a very large causal impact on their teams' performance, AFL teams appear to underpay superstar players by at least 30 per cent. We offer a number of potential explanations for this behaviour. We also note that while prior economics research has examined superstar productivity and pay in professional sports, our paper is the first study that we are aware of that explores these concepts in a setting – the AFL – that employs a 'hard' salary cap. We exploit this feature of our setting to discuss the role that labour market regulations may play in dampening superstar pay.[9] Finally, by presenting evidence that suggests that superstar effects are isolated to players in the 95th percentile, we also provide an empirically grounded definition of what constitutes a superstar player in the context of the AFL.

The rest of the paper reads as follows. Section II describes the paper's setting and the data used. Section III explains the empirical strategy used. In Section IV, we report and interpret our results. Section V contains a number of robustness checks that evaluate the plausibility of our empirical strategy. In Section VI, we conclude.

II Setting and Data

The AFL is the nation's premier Australian-rules football competition. Australian-rules football is the most popular sport in Australia, and the AFL is by the far the country's most commercially successful and well-supported sports competition.[10] In this paper, we use game-level data on all regular season AFL games played between

[8] Interested in how differences in club finances affect competitive balance in the AFL, Borland and Booth (2014) show that only in more recent seasons are teams' football operations expenditures related to their on-field success.

[9] Our discussion of the role that the salary cap plays in our finding that superstars are potentially underpaid by 30 per cent has implications for the ongoing debate over whether AFL clubs should be permitted to pay a marquee player outside of the salary cap (SEN, 2020). Several professional sports leagues employ a marquee player rule, including the AFLW, where teams are permitted to pay two players a wage inside the salary cap and an additional amount outside of the cap. If the AFL adopted a marquee player rule, we may well observe that superstar players are paid in line with their marginal product.

[10] As per its annual report, the AFL generated A$668 million in revenue during the 2018 season. For comparison, the National Rugby League – the second-largest sports league in Australia – reported A$500 million in revenue for the 2018 season.

2013 and 2018. This dataset contains information on game outcomes, team line-up changes, player performance rankings, and betting odds. The dataset was constructed from three different sources: official AFL match reports, performance data generated by the AFL's data provider, Champion Data, and online bookmaking sites. Given our study exploits game-to-game line-up changes, we drop from our sample all games from the first week of each season (by definition, there are no line-up 'changes' at the start of each season). As such, our dataset has a total of 1133 game-level observations.

We use Champion Data's player-game level ranking points to define superstar players. During each game, Champion Data rates player performance using a single metric that is tied to a large number of quantitative performance measures (e.g., effective disposals, goals, contested marks, meters gained, etc). Champion Data's player points are the league's most-widely used and reported quantitative performance rating. The objective of this rating is to capture a player's overall contribution to his team's performance.[11] For the purposes of this study, we calculate AFL players' average game-level performance rating for each season in our sample. We then classify player i as a superstar in season t if his average performance rating in season $t - 1$ was in the 95th percentile of the league's performance distribution for that season.[12] If superstar talent is uniformly distributed across the competition (a reasonable assumption given the AFL's use of reverse-order drafts and hard salary caps), by our definition the average AFL team's list has approximately two superstars (the typical AFL team has a list of 40 players, and we assume superstars make up 5 per cent of players in the competition).[13]

To estimate the marginal productivity of superstar players, we exploit teams' game-to-game injury-induced line-up changes. AFL policy requires clubs to announce their team line-up either one or two days prior to an upcoming game. As part of these line-up announcements, teams identify the 22 players that they have selected to play in that week's game.[14] When line-ups are announced, teams explicitly identify two sets of players: 'ins' and 'outs'. 'Ins' are players that did not feature in the previous week's line-up and have been brought into the team for the upcoming week's game (e.g., players that are returning from injury, or have been promoted from the club's reserve team). 'Outs' are players that featured in the previous week's line-up but are not in the team named for the upcoming week's game. Typically, these are players that have been injured in the previous week's game or during mid-week training in preparation for the upcoming week's game. These players may also be uninjured athletes who have been dropped due to poor form or players that the team has chosen to 'rest' for the upcoming game.

In addition to explicitly identifying each of their 'outs', teams must also disclose the reason for each line-up change. If a player is injured, the specific injury or illness is disclosed. If a player is dropped or 'rested', he is listed as 'omitted' or

[11] For more information on the player ranking points, see https://www.championdata.com/glossary/afl/.

[12] In our sample, average game-level performance ratings between season $t - 1$ and t are strongly positively correlated ($\rho = 0.66$ and $P < 0.01$). We also find that more than half of the players in the 95th percentile in season $t - 1$ are also in the 95th percentile in season t. This suggests that while our measure is persistent, it successfully captures a realistic amount of churn in superstar status, that is, across seasons we expect that young ascendant talents to squeeze out fading veterans for superstar status, while a large core of superstars in their prime maintain their status.

[13] We find robust empirical support for our assumption that the average list has two superstars. In our sample, as we define them, superstars are evenly distributed across the competition each season, and while high-performing teams, on average, do have more superstars on their lists, it is very uncommon for a team to not have any superstars in a given season. For example, we find that across the period 2013–18, on average, approximately 17 of the AFL's 18 teams have at least one superstar on their playing list, while all teams have at least one superstar in one of the seasons that we examine. Furthermore, conditional on a team having at least one superstar, we find that the average team each season has 2.2 superstars. These findings suggest that the actual distribution of superstars is not highly skewed to the best performing teams.

[14] Teams also declare a set of 'emergency' players, who can be brought into the line-up in the event that any of the starting 22 players are required to pull out prior to the game. Before 2018, teams named three emergency players; from 2018, teams have been required to name four emergency players.

'rested' in the line-up announcement. As such, in our analysis we isolate players who teams were unable select (due to injury) from players that teams 'chose' not to select (due to form, disciplinary reasons, or list management purposes, e.g., rest). For the purposes of identification, we only exploit the former source of variation.

In our dataset, we also observe the 'type' of injury sustained by each player in our sample. While some teams rest players under the guise of injuries (often citing 'soreness', 'unwell', or 'managed' as the reason for the line-up change), we note that none of the superstar injuries in our sample are classified as such. Instead, we observe that superstars in our sample are only subject to 'acute' injuries, for example, muscle tears, joint damage, concussions or bone fractures. We believe that this feature of our sample at least partially alleviates the concern that at certain stages of the season or in advance of games against particular opponents teams may endogenously rest superstars under the guise of injury (for instance, because the team was already a near certainty to win or lose the upcoming game, or because the result of the upcoming game had little bearing on the final standings for the season).[15] By construction, there is a direct mapping between 'outs' and 'ins' (the latter replace the former).

While injuries to superstars are not common, they do occur with some frequency (we discuss the frequency and distribution of superstar injuries in Section IV). When injury forces a team to remove a superstar from its line-up, this player is replaced by a player on the team's list who is not currently in the line-up (i.e., a player that was not selected for the previous game). We refer to such players as replacement-level players. As each team has a list of approximately 40 players, and must select a line-up of 22 players each week, teams in theory have approximately 18 players to select from when replacing an injured player. However, not all these players are 'on the margin'. As a share of these players will be rookies that have been recently drafted and may not be ready for league football, and another share will be players on the long-term injury or disabled list (who are unable to play), teams usually have eight to 10 replacement-level players to select from when covering for an injured starting player. In this role, these replacement-level players – who we can think of as 'on the margin' for selection each week – play approximately six to eight games per season.[16] We use the details of these announcements and players' performance ratings to identify instances where due to injury superstar players were removed from their team's line-up between games.

Finally, in our analysis, we also make use of game-level betting market data. In our setting, weekly team line-up announcements are closely followed by the public (line-ups are discussed intensely on local sports television programs and on the Internet) and by bookmakers (sports betting in Australia is legal and betting on the outcome of individual AFL games is extremely active). In general, bookmakers set their 'opening odds' for an upcoming game at the start of each week. Over the course of the week, bookmakers will adjust these odds in response to market forces and the disclosure of information material to the outcome of the upcoming game (e.g., weather forecasts, interviews and press

[15] Under the pretence of 'injuries', in-form clubs might strategically rest their superstars toward the end of the season so as to 'freshen up' them for the first round of the finals. For example, in 2015, Fremantle famously rested 11 players (half the previous week's line-up) for an inconsequential end-of-season game. Similarly, clubs toward the bottom of the ladder that face a challenging string of matches toward the end of the season might send superstar players off for surgery before the season has finished so that these players can recover in time to train during the following pre-season. In both these scenarios, superstar injuries are endogenous to the upcoming game's outcome (in the former scenario, such endogeneity will likely bias our superstar estimates downwards, that is, superstars are non-randomly missing from easier-to-win games; while in the latter scenario, such endogeneity will likely bias our superstar estimates upwards, that is, superstars are non-randomly missing from harder-to-win games). While a concern in theory, notes 23 and 27 below present empirical evidence that suggests that our identification strategy does not appear to be violated by 'strategic' injury-induced line-up changes and that superstar injuries are plausibly exogenous.

[16] A regular team player who has recovered from injury may come back into the line-up when injury causes a superstar to withdraw. However, our 'control group' in this instance is still a replacement-level player as presumably a replacement-level player came into the line-up initially in place of the injured regular team player, that is, the superstar's injury keeps the replacement-level player in the line-up, which for the purposes of our empirical strategy is the same as the replacement-level player coming into the team line-up to cover for the superstar.

conferences with players and coaches, etc). In this way, the betting odds come to impound the information contained in line-up announcements. For this reason, the disclosure of injuries to superstar players should lead to large changes in the betting odds for a specific game. We argue that these changes reflect shifts in the public's expectations about whether home team i will win game t. Bookmakers continue to take bets and adjust the odds up until the start of each game, at which point the 'closing' odds are set. To study how betting markets price superstars' marginal productivity, in our analysis we examine variation in bookmakers' closing odds and closing line (i.e., the 'spread').

III *Empirical Strategy*

(i) Identification

For injury-induced line-up changes to identify superstar players' marginal productivity, injuries need to be exogenous. In other words, superstars' injuries must only relate to teams' future performance outcomes through direct effects that arise from teams having to drop these players from their line-up. In practice, this means that injury-induced line-ups changes between games $t - 1$ and t cannot be endogenous to: the quality of the teams competing in game t; when game t takes place during the season; the weather and playing conditions of game t; or, any other factor related to the competing teams' observed performance outcomes in game t. While this assumption cannot be empirically tested, based on key institutional details and features of our setting, we argue that injuries to superstar players in the AFL are indeed plausibly exogenous.

First, we emphasise that Australian-rules football is an especially physically demanding sport. According to Hrysomallis (2013), the most common injuries suffered by AFL players are blunt force trauma (e.g., head or body knocks from tackles and collisions) and joint injuries (e.g., torn knee, shoulder and ankle ligaments). These injuries do not occur in any sort of systematic manner, but usually arise due to random on-field events (e.g., a player slips over during a tackle and strikes his head, a player injures his knee sliding after a ball in the wet conditions). Even if we allow that injuries are more likely to occur in closer games or in especially important contests when players might take greater physical risks or exert themselves excessively, injury-induced line-up changes between games are still likely to be exogenous in that our empirical strategy exploits injuries in games $t - 1$ as negative shocks on a team's line-up for game t, that is, we are not examining the relationship between superstar injuries in game t and team performance in game t. To follow up this feature of our setting, we also emphasise that the AFL does not order the fixture based on team quality, for example, a club that faces a good team in game t is not more likely to face a weak team in game $t + 1$. This reduces the concern that game t's attributes are correlated with game $t + 1$'s attributes and lends further credence to our argument that injury-induced line-up changes are plausibly exogenous (in Section IV.ii, we provide further empirical evidence to support this conclusion).[17]

However, our identification strategy does acknowledge that specific teams are likely to be better (or worse) at managing injuries. This could arise because more successful teams have better-resourced sports science and medical departments, which in turn allows these teams to better manage and prevent injuries (McCall et al., 2014). We also acknowledge that teams may employ strategies, tactics or playing styles that are correlated with injury rates and on-field performance. For instance, teams that move the ball quickly by foot may win more games but may also suffer greater rates of high-impact trauma injuries. For these reasons, in our empirical design, we use team fixed effects. Formally,

[17] While not likely given the risk involved, we acknowledge that in certain circumstances teams may use their discretion to keep an injured superstar in the line-up. This may arise if the superstar has a marginal injury and if the upcoming game is especially important, for example, by winning the game the team secures a place in the finals. While in our main results and robustness checks we present evidence that suggests that a team's decision to omit or retain an injured player is not endogenous, we also emphasise that even if superstar injuries are endogenous in this manner, such a violation should act as a downward bias on our estimate of superstars' marginal productivity, that is, as injuries likely lower superstars' performance, then average team performance in our 'control' group should be dampened relative to average team performance in the potential outcome/unobserved counterfactual we would ideally use for identification. If we allow for this source of potential endogeneity, our paper offers a lower bound estimate of superstars' marginal productivity.

this means that we exploit within-team variation in injury-induced line-up changes and performance to identify superstar players' marginal productivity. For this approach to address the endogeneity concerns raised above, we must also assume that team finances and playing styles are 'sticky' over the short to medium term (i.e., that team finances and performance are time-invariant over the period of our study).[18]

(ii) Estimation

In this paper, the outcome measure of interest is win_{ijs}, a binary variable that is equal to 1 for games where home team i defeats away team j in season s, and 0 otherwise. We also examine a continuous outcome measure, $\text{margin}_{\text{ijs}}$. This variable is the final point spread or margin of the game between home team i and away team j in season s (when $\text{margin}_{\text{ijs}} > 0$, home team i wins game ijs, and when $\text{margin}_{\text{ijs}} < 0$, home team i loses game ijs). As our analysis focuses on game outcomes (or expected game outcomes), our estimates do not capture the value that teams extract from superstars' marketing and off-field productivity (e.g., revenue from sponsorship deals and increased attendance, membership and merchandise sales). As such, we present our empirical findings as lower-bound estimates of superstars' total productivity.[19] We sourced all data on game outcomes and attributes from official AFL match reports.

To estimate the marginal productivity of superstar players, we use OLS to estimate the following set of reduced-form equations:

$$\begin{aligned}\text{win}_{\text{ijs}} = {} & \alpha_0 + \alpha_1 \text{home star injuries}_{\text{ijs}} \\ & + \alpha_2 \text{away star injuries}_{\text{ijs}} + \alpha X_{\text{ijs}} \\ & + \phi_i + \chi_s + \epsilon_{\text{ijs}}\end{aligned} \tag{1}$$

$$\begin{aligned}\text{margin}_{\text{ijs}} = {} & \beta_0 + \beta_1 \text{home star injuries}_{\text{ijs}} \\ & + \beta_2 \text{away star injuries}_{\text{ijs}} \\ & + \beta \text{X}_{\text{ijs}} + \tau_i + \gamma_S + \eta_{\text{ijs}}\end{aligned} \tag{2}$$

We define $\text{home star injuries}_{\text{ijs}}$ as the number of star players who, due to injury, were removed from home team i's line-up for the game against away team j in season s. We define $\text{away star injuries}_{\text{ijs}}$ as the number of star players who, due to injury, were removed from away team j's line-up for the game against home team i in season s. By allowing these measures to separately enter Equations (1) and (2), we let the effect of superstars on team performance vary across home and away teams (in Section IV.iii, we offer explanations for why we expect this to be the case). X_{ijs} is a vector of game-level covariates, ϕ_i and τ_i are home-team fixed effects, χ_s and γ_s are season fixed effects, and ε_{ijs} and η_{ijs} are idiosyncratic error terms.[20]

α_1 and α_2, and β_1 and β_2 are our estimates of superstar players' marginal productivity. α_1 is the effect of home team i's superstar player on the likelihood that home team i wins game ijs. α_2 is the effect of away team j's superstar player on the likelihood that home team i wins game ijs.[21] β_1 is the effect of home team i's superstar player on

[18] In Section IV, we relax this assumption and report results from a series of robustness checks that use team-season fixed effects. Under these alternative specifications, we allow team finances and playing style to vary from season to season, that is, we exploit within-team-season variation. Our results do not change when we employ this approach.

[19] For our estimates to serve as a lower bound, we must assume that superstars' off-field productivity is non-negative, that is, that superstars, through their marketing and off-field activities, do not, on average, diminish team revenue. We believe that this assumption is not unreasonable, and we argue that it is consistent with historical recruitment practices in the AFL where expansions teams in new markets signed superstar players to attract popular attention and boost membership numbers, for example, the Brisbane Bears signing Warwick Capper and the Gold Coast Suns recruiting Gary Ablett, Jr.

[20] As we mention in the introduction, we use specifications that are linear in the number of injured superstars. As we only observe very few games where teams lose two superstars to injury, we cannot separately estimate the effects of losing a second or third superstar on team performance.

[21] As stated above, we use ordinary least squares (OLS) to estimate α_1 and α_2. Linear probability models, that is, OLS regressions on binary variables, are convenient, computationally tractable and often likely to have less bias than index model alternatives such as probit and logit (Horrace & Oaxaca, 2006; Angrist & Pischke, 2008; Wooldridge, 2010). Due to the incidental parameters problem, linear probability models are also better suited to the use of fixed effects (Greene, 2004). Nonetheless, to assuage the concern that specification choice is driving our results, in untabulated analysis (available from the authors upon request), we show that logit and probit models produce estimates of α_1 and α_2 that are qualitatively similar to our main results.

TABLE 1
Descriptive Statistics

	N	Mean	*SD*	Minimum	Maximum
Endogenous variables					
Win	1133	0.56	0.50	0	1
Margin	1133	5.94	44.98	−138	148
Prob. Home win	1133	0.55	0.26	0.02	0.98
Expected margin	1131	6.45	29.33	−87.50	101.50
Exogenous variables					
Home Star Injuries (95th)	1133	0.02	0.13	0	1
Away star injuries (95th)	1133	0.02	0.16	0	2
Home star injuries (90th)	1133	0.06	0.24	0	2
Away star injuries (90th)	1133	0.08	0.30	0	2
Home star injuries (90–95th)	1133	0.04	0.20	0	1
Away star injuries (90–95th)	1133	0.06	0.25	0	2
Home standings	1124	9.48	5.23	1.00	18.00
Away standings	1124	9.50	5.17	1.00	18.00
Controls					
Home Elo rating	1133	1499.85	143.72	1214.18	1796.13
Away Elo rating	1133	1499.93	143.50	1213.37	1798.35

Notes: Descriptive statistics for the variables used are shown. win_{ijs} is a binary variable that is equal to 1 for games where home team *i* defeats away team *j* in season *s*, 0 otherwise. $\text{margin}_{\text{ijs}}$ is the final point spread or margin of the game between home team *i* and away team *j* in season s. prob home win_{ijs} is the betting odds-implied probability that home team *i* defeats away team *j* in season *s*. This measure uses the average bookkeepers' margin-adjusted closing odds. expected $\text{margin}_{\text{ijs}}$ is the expected margin of the game between home team *i* and away team *j* in season *s*. This measure is based on the average bookkeepers' line or points spread. home star $\text{injuries}_{\text{ijs}}$ is the number of star players who, due to injury, were removed from home team *i*'s line-up for the game against away team *j* in season *s*. away star $\text{injuries}_{\text{ijs}}$ is the number of star players who, due to injury, were removed from away team *j*'s line-up for the game against home team *i* in season *s*. Depending on the specification, we classify player *i* as a superstar in season *t* if his average performance rating in season *t* – 1 was in either the 95th, 90th, or 90th–95th percentiles of the league's performance distribution. home $\text{standing}_{\text{ijs}}$ is home team *i*'s position in the standings as of their game against away team *j* in season *s*. away $\text{standing}_{\text{ijs}}$ is away team *j*'s position in the standings as of their game against home team *i* in season *s*. home elo $\text{rating}_{\text{ijs}}$ is home team *i*'s Elo rating as of their game against away team *j* in season *s*. And away elo $\text{rating}_{\text{ijs}}$ is away team *j*'s Elo rating as of their game against home team *i* in season *s*.

home team *i*'s margin in game ijs. β_2 is the effect of away team *j*'s superstar player on home team *i*'s margin in game ijs. As we define win_{ijs} and $\text{margin}_{\text{ijs}}$ from home team *i*'s perspective, we expect $\alpha_1 < 0$ and $\alpha_2 > 0$, and $\beta_1 < 0$ and $\beta_2 > 0$.

We emphasise that in our empirical strategy α_1, α_2, β_1 and β_2 are identified at the margin, that is, these coefficients are estimates of superstar players' productivity over and above replacement-level players (e.g., the players who on average step in to replace injured superstars). For our empirical strategy to identify the marginal productivity of superstar players, injury-induced line-up changes must be uncorrelated with ε_{ijs} and η_{ijs}. In Section III.i, we discussed features of our setting that lend support to this assumption. While this assumption cannot be formally tested, in Section IV, we provide empirical evidence consistent with the argument that injury-induced line-changes are plausibly exogenous.

IV Results

(i) Descriptive Statistics

In Table 1, we present summary statistics. As we noted above, our sample contains information on 1133 AFL games played between 2013 and 2018.[22] In terms of our outcome measures, we find that the home team wins 56 per cent of games and we obverse that on average the home team wins by six points (or one goal). We note that $\text{margin}_{\text{ijs}}$ displays considerable variation (its standard deviation is more than seven times the size of its mean) and that our sample contains games where the home team won by as many as 148 points and games where the home team lost by as many as 138 points.

[22] For *expected* $\text{margin}_{\text{ijs}}$, we are missing observations for a small subset of these games. This is because online betting data about the expected margin for these games do not appear to be available.

TABLE 2
Exogeneity of Star Injuries: Regressions of Star Injuries on Game and Team Attributes

	Dependent variable					
	Home star injuries (1)	Away star injuries (2)	Home star injuries (3)	Away star injuries (4)	Home star injuries (5)	Away star injuries (6)
Home standings	0.001	0.001				
	(0.001)	(0.001)				
Away standings	−0.001	0.0003				
	(0.001)	(0.001)				
Opening odds for home Win			0.003**	−0.001		
			(0.001)	(0.002)		
Home Elo					−0.00001	−0.0001
					(0.00003)	(0.00003)
Away Elo					0.00004	−0.0001
					(0.00003)	(0.00004)
Yearly trend	−0.003	−0.003	−0.002	0.003	−0.003	−0.003
	(0.002)	(0.003)	(0.002)	(0.003)	(0.002)	(0.003)
Round trend	0.001	0.001	0.0005	0.001	0.0005	0.001
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Thursday–Friday	0.001	−0.028	0.001	−0.027	−0.0005	−0.022
	(0.035)	(0.043)	(0.035)	(0.043)	(0.035)	(0.043)
Saturday	0.018	−0.041	0.019	−0.043	0.021	−0.046
	(0.033)	(0.040)	(0.033)	(0.041)	(0.033)	(0.041)
Sunday	0.016	−0.039	0.017	−0.041	0.018	−0.044
	(0.033)	(0.041)	(0.033)	(0.041)	(0.033)	(0.041)
Evening slot	0.015	−0.017	0.016	−0.017	0.014	−0.015
	(0.011)	(0.013)	(0.011)	(0.013)	(0.011)	(0.014)
Observations	1124	1124	1133	1133	1133	1133
R^2	0.007	0.005	0.009	0.005	0.007	0.009

Notes: *$P < 0.1$; **$P < 0.05$; ***$P < 0.01$. Estimates from ordinary least squares (OLS) regressions of our superstar injury measures on a range of observable game and team characteristics are presented. Table 1 defines each variable examined in this paper. Standard errors are reported in parentheses below the coefficient estimates.

In terms of our betting variables, we find that the market expects the home team to win 55 per cent of games. Given that the average values of win_{ijs} and prob win_{ijs} are so similar, we conclude that the AFL betting market appears to be highly efficient. In support of this point, we also note the similarity in the average values of $\text{margin}_{\text{ijs}}$ and expected $\text{margin}_{\text{ijs}}$. Again, we take this as evidence that betting markets do a reasonable job of pricing (public) information that is material to the outcome of an upcoming game.

Finally, in terms of our exogenous variables, we observe that, while not common, injury-induced line-up changes that involve superstar players affect a sizeable portion of games in our sample. We note that, while very rare, there are a small number of instances where injuries forced the away team to withdraw multiple superstars from its line-up.

(ii) Injuries as Plausibly Exogenous Shocks

As we discussed in Section III, for injury-induced line-up changes to identify superstars' marginal productivity, injuries must be plausibly exogenous. While we cannot directly test this, in Table 2, we present evidence that is broadly consistent with our identifying assumption. Specifically, we report the results from multivariate analyses where we regress home star $\text{injuries}_{\text{ijs}}$ and away star $\text{injuries}_{\text{ijs}}$ on a range of observable 'pretreatment' characteristics of the teams competing in game ijs.

Across each of the columns in Table 2, we show that home star $\text{injuries}_{\text{ijs}}$ and away star $\text{injuries}_{\text{ijs}}$ do not appear to be systematically related to time-varying factors. That is, our estimates suggest that superstar injuries are not increasing or decreasing season on season. Superstar injuries also do not appear to be increasing or decreasing over time within each season.[23] Across each of the columns in Table 2, we also show that home star $injuries_{ijs}$ and away star $injuries_{ijs}$ do not appear to be systematically related to factors tied to league schedule or fixturing. That is, our estimates suggest that superstar injuries are no more or less common when a game is played on a specific day of the week (on average, games played earlier in the week are preceded by a shorter rest period for the players, which in theory could be associated with injury-induced line-up changes). Superstar injuries also do not appear to be related to whether a game is played in the afternoon or the evening.

In Table 2, we also examine how injury-induced line-up changes are related to a range of measures of team quality. In columns (1) and (2), we show that superstar injuries do not appear to be related to the position in the standings of the competing teams. In columns (3) and (4), we examine whether opening odds for a home team win are related to superstar injuries. While we find insignificant results for superstar injuries to the away team, we do find that superstar injuries to the home team are positively related to the opening odds for a home team win (longer odds imply that a home team win is less likely). While one interpretation of this result is that home teams may be more likely to rest superstar players under the guise of injury for games that the team are expected to lose, another explanation is that there may be 'leakage' in the disclosure of injury-induced line-up changes. For example, if a superstar player is clearly badly injured during game $t-1$, the market may correctly infer prior to the actual mid-week line-up announcement that the player will not appear in game t. As the opening odds are set before the mid-week line-up announcement, this means these types of injury shocks will already be priced into the opening odds, therefore giving the impression that injury induced line-up changes are more likely to occur in advance of games that the home team is expected to lose.[24] In columns (5) and (6), we show that superstar injuries to the home or away team are unrelated to home and away team quality as measured by Elo ratings.[25] Across columns (1) to (6), our results appear to consistently suggest that home star $injuries_{ijs}$ and away star $injuries_{ijs}$ are not related to the quality of teams competing in game ijs.

Overall, by showing that superstar injuries appear to be unrelated to a range of observable pre-treatment outcomes, Table 2 lends support to the idea that injuries are plausibly exogenous.

(iii) Main Results

In Table 3, we present our estimates of superstar players' marginal productivity. Specifically, in columns (1) to (4), we report our results from regressions of win_{ijs} on home star $injuries_{ijs}$ and away star $injuries_{ijs}$. In columns (5) to (8), we report our results from regressions of $margin_{ijs}$ on home star $injuries_{ijs}$ and away star $injuries_{ijs}$. Across columns (1) to (8), the estimated coefficients on home star $injuries_{ijs}$ and away star $injuries_{ijs}$ do not appear to vary greatly when we add home team, season, and home team-season fixed effects, nor when we also control for home and away team quality using the teams' position in the standings as of game ijs.

If we use column (4) as our preferred specification of Equation (1), we find that home star $injuries_{ijs}$ are not related to win_{ijs} at any conventional level of statistical significance. We note, however, that, as predicted in Section III.ii, the sign of the coefficient is negative (albeit not statistically distinguishable from zero when using a two-tailed test). We do observe that away star $injuries_{ijs}$ are positively related to win_{ijs} at the 5 per cent level. In terms of economic significance,

[23] We acknowledge that a linear time trend to address within-season temporal variation in injuries involves somewhat restrictive functional form assumptions. As such, in an untabulated analysis we show that superstar injuries are no more likely to occur in games that are more likely to be 'dead rubbers' (i.e., games where the result does not have a bearing on the final standings for the season and so teams may have incentives to 'rest' star players and therefore avoid the risk of injuries before the finals). Specifically, we re-run the specifications in Table 2 while also including an indicator variable that is equal to 1 for games that take place in the last four rounds of the season. We argue that at this stage of the season teams have likely already qualified or failed to qualify for the finals.

[24] In this sense, opening odds are perhaps not best classified as a 'pre-treatment' characteristic in our research design. Of course, another more benign explanation for the significant coefficient on opening odds for a home team win is that we are simply picking up a random association between injuries and odds, in which case this association has no 'economic meaning'.

[25] First developed to rank players in chess, Elo ratings are commonly used in games and sport to measure the quality of participants. A contestant's Elo rating is represented by a number that increases or decreases depending on the outcome of games between rated contestants.

TABLE 3
Measuring the Productivity of Star Players: Regressions of Team Performance on Star Injuries

	Dependent variable							
	Win				Margin			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Home star injuries	−0.172	−0.161	−0.082	−0.148	−13.856	−1.204*	−3.699	−6.041
	(0.147)	(0.118)	(0.104)	(0.095)	(8.696)	(5.912)	(5.670)	(7.136)
Away star injuries	0.138**	0.151***	0.139***	0.168**	12.703*	15.243***	14.124***	15.975**
	(0.056)	(0.044)	(0.043)	(0.068)	(7.103)	(5.126)	(4.587)	(6.476)
Home standing			−0.042***				−3.867***	
			(0.002)				(0.262)	
Away standing			0.036***				3.651***	
			(0.002)				(0.163)	
Home team FEs	No	Yes	Yes	No	No	Yes	Yes	No
Season FEs	No	Yes	Yes	No	No	Yes	Yes	No
Home team-season FEs	No	No	No	Yes	No	No	No	Yes
Observations	1133	1133	1124	1133	1133	1133	1124	1133
R^2	0.004	0.098	0.402	0.233	0.004	0.119	0.465	0.305

Notes: $^*P < 0.1$; $^{**}P < 0.05$; $^{***}P < 0.01$. Estimates from ordinary least squares (OLS) regressions of game outcomes on our superstar injury measures are presented. Table 1 defines each variable examined in this paper. Standard errors, reported in parentheses below the coefficient estimates, are clustered at either the home-team or home team-season level.

the coefficient on away star injuries$_{ijs}$ in column (4) implies that, for away teams, a superstar – compared with a replacement-level player – increases his team's chance of winning by approximately 17 percentage points.[26]

If we use column (8) as our preferred specification of Equation (2), we find that home star injuries$_{ijs}$ are not related to margin$_{ijs}$ at any conventional level of statistical significance. Again, we note, however, that, as predicted in Section III.ii, the sign of the coefficient is negative (albeit not statistically distinguishable from zero when using a two tailed test) We do observe that away star injuries$_{ijs}$ are positively related to margin$_{ijs}$ at the 5 per cent level. In terms of economic significance, the coefficient on away star injuries$_{ijs}$ in column (8) implies that, for away teams, a superstar – compared with replacement-level players – improves his team's final margin in game ijs by approximately +16 points (or, equivalently, 2.6 goals).[27]

We suggest several reasons for why away star injuries$_{ijs}$ but not home star injuries$_{ijs}$ are significantly related to home team *i*'s performance outcomes. First, superstar effects may exist in both home and away games, but our empirical strategy may not be sufficiently powered to detect the former. We note that our sample contains 25 games where

[26] Speaking to the economic magnitude of this effect, we note that an AFL team's line-up consists of 22 players. We also note that the home team wins 56 per cent of games in our sample. As such, our estimates suggest that even as just a single player on the field a superstar in the AFL exerts a very large impact on a game's eventual outcome.

[27] In untabulated analysis, we also re-run the specifications in Table 3 while including a 'dead rubber' variable – an indicator that is equal to 1 for games that take place in the final four rounds of the season. We do so to address the concern that superstar injuries may be endogenous to games that occur late in the season when the final standings are already largely decided, that is, when teams may opt to 'rest' players because the regular season is 'over'. However, we find our main results remain unchanged when we also control for these 'dead rubbers'. Similarly, our results remain largely the same when we re-estimate the specifications in Table 3 on the subset of games that occur in the first three-quarters of the home and away season, that is, when the overwhelming majority of games still have some bearing on the final standings for the season. We take this as further evidence that superstar injury-induced line-up changes in our setting are plausibly exogenous.

superstar injuries affected the away team. In contrast, our sample contains 20 games where superstar injuries affected the home team.[28] This issue of statistical power may explain why in Table 3 our estimates of superstar effects in home games are of a similar magnitude but less precise than our estimates of superstar effects in away games.

If we put to one side this issue of statistical power, another explanation for this pattern of results is that away games are harder to win and that superstars' marginal productivity only emerges under the more challenging conditions of playing away from home.[29] An example of why this might occur is if superstar players are more effective than replacement-level players at neutralising the opposing team's home-ground advantage.[30] This could happen if superstars are better able to adapt their playing style or strategy to the unique dimensions of the away teams' playing ground, that is, replacement level players have highly 'firm' specific human capital, whereas superstars have more general human capital. This could also happen if away-team superstars are able to reverse or neutralise officials' home bias (i.e., due to increased public scrutiny umpires may be less likely to unfairly penalise an away team player if that player is a high profile star) or if away-team superstars are less likely to under-perform or choke under pressure from a hostile home crowd (in this sense, superstars are 'clutch' performers). A final explanation for our results is that superstars may be subject to lower 'travel costs'. According to this view, superstars outperform replacement-level players because the former, due to better preparation or greater inherent resilience, are less affected by inter-state travel or less disrupted by 'living' on the road (e.g., reduced air quality and oxygen pressure from extended air travel may cause respiratory and circulatory issues, living out of hotel accommodation may affect sleep and nutrition, etc).[31]

In Table 4, we present betting market-based measures of superstar players' marginal productivity. Specifically, in columns (1) to (4), we report our results from regressions of prob win_{ijs} on home star $injuries_{ijs}$ and away star $injuries_{ijs}$. In columns (5) to (8), we report our results from regressions of expected $margin_{ijs}$ on home star $injuries_{ijs}$ and away star $injuries_{ijs}$. Across columns (1) to (8), the estimated coefficients on home star $injuries_{ijs}$ and away star $injuries_{ijs}$ do not appear to vary greatly when we add home team, season, and home team-season fixed effects. When we control for home and

[28] Reflecting these differences in power, for all specifications in Table 3, the standard errors on the estimate for home star $injuries_{ijs}$ are larger than the standard errors on the estimate for away star $injuries_{ijs}$. In columns (1) to (3), the former are approximately twice the magnitude of the latter.

[29] To shed further light on this potential explanation, in untabulated analysis (available from the authors upon request) we examine player–game-level performance data and compare superstars' and non-superstars' player performance ratings across home and away games for the period 2013–18. We find that, on average, superstars have much higher performance ratings than non-superstars, and that the magnitude of this performance differential is largely the same across both home and away games. On face value, this result suggests that teams do not appear to adopt specific strategies for away games that rely more heavily on superstar players, and that issues of statistical power provide a more compelling explanation for our pattern of findings in Table 3. However, we stress that these results should be interpreted with caution as players' performance is not perfectly observable and that superstars may in fact generate excess wins over replacement-level players in away games, but do so by elevating their output along performance dimensions that are not well captured by the AFL performance rating, for example, superstars in away games perform better in 'the clutch' (e.g., win pivotal free kicks or score high-pressure goals) or increase their workload 'off the ball'.

[30] Consistent with such a home advantage in our setting, we show in Table 1 that for our sample the home team wins 56 per cent of games. We also note that betting odds imply that the home team is expected to win 55 per cent of games. More generally, a large literature in sports economics examines the magnitude and underlying mechanism of home-ground advantage in professional sports (Stefani & Clarke, 1992; Clarke & Norman, 1995).

[31] Given that home games are 'easier' to win, there is also a potentially straightforward mechanical reason for why superstar effects may only show up in away games. At the extreme, in a game where the home team is a very strong favourite, the marginal impact of a home-team superstar on the probability that the home team wins must approach zero (i.e., absent the superstar, the home team nonetheless remains a near certainty to win). If our sample is composed of a sufficient number of these extreme games, we will fail to find home-team superstar effects. An issue with this explanation is that, if injury-induced line-up changes occur at random, the inverse should also hold: at the extreme, in a game where the home team is a very strong favourite, the marginal impact of an away-team superstar on the probability that the home team wins must approach zero (i.e., absent the superstar, the away team nonetheless remains a near certainty to lose). As we show that away-team superstar effects are economically and statistically significant, we fail to observe that this is the case.

TABLE 4

Measuring the Productivity of Star Players: Regressions of Betting Market Expectations on Star Injuries

	Dependent variable							
	Prob. home Win				Expected margin			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Home star injuries	−0.144**	−0.124**	−0.068**	−0.107**	−16.648**	−14.478*	−8.053**	−12.577**
	(0.063)	(0.060)	(0.027)	(0.053)	(7.393)	(7.475)	(3.559)	(6.186)
Away star injuries	0.050	0.067	0.056	0.059	7.235	9.181*	7.742*	8.215
	(0.052)	(0.043)	(0.036)	(0.047)	(5.914)	(4.980)	(4.072)	(5.315)
Home standing			−0.022***				−2.440***	
			(0.002)				(0.295)	
Away standing			0.028***				3.207***	
			(0.001)				(0.084)	
Home team FEs	No	Yes	Yes	No	No	Yes	Yes	No
Season FEs	No	Yes	Yes	No	No	Yes	Yes	No
Home team-season FEs	No	No	No	Yes	No	No	No	Yes
Observations	1133	1133	1124	1133	1131	1131	1122	1131
R^2	0.006	0.232	0.711	0.451	0.007	0.220	0.709	0.444

Notes: $^*P < 0.1$; $^{**}P < 0.05$; $^{***}P < 0.01$. Estimates from ordinary least squares (OLS) regressions of betting market outcomes on our superstar injury measures are presented. Table 1 defines each variable examined in this paper. Standard errors, reported in parentheses below the coefficient estimates, are clustered at either the home-team or home team-season level.

away team quality using the teams' position in the standings as of game ijs, our estimates are somewhat attenuated but nonetheless remain economically and statistically significant.

Across columns (1) to (8), we find weak statistical evidence that, as predicted in Section III.ii, the coefficient on away star injuries$_{ijs}$ is positive. While only two of our specifications produce coefficients on away star injuries$_{ijs}$ that are statistically significant at conventional levels when using two-tailed tests, we find that when using one-tailed tests, for all specifications in Table 4, the evidence supports the alternative hypothesis that the coefficient on away star injuries$_{ijs}$ is greater than zero at at least the 10 per cent level.[32]

[32] The existence of superstars is a widely accepted position in professional sports. As such, we have a strong prior belief regarding the direction of the coefficients on our superstar injury variables, that is, superstars are more productive than replacement-level players (this is not the same as saying we have a strong prior about the degree to which superstars are more productive than replacement-level players, this paper's primary motivating question). Given this prior, we feel that our one-tailed tests provide at least some evidence that the coefficient on away star injuries$_{ijs}$ is positive.

Next, we observe that home star injuries$_{ijs}$ are positively related to prob win$_{ijs}$ at either the 5 per cent or 10 per cent level. In terms of economic significance, the coefficient on home star injuries$_{ijs}$ in column (4) implies that, for home teams, a superstar – compared with a replacement-level player – increases his team's expected chance of winning by approximately 11 percentage points. In terms of economic significance, the coefficient on home star injuries$_{ijs}$ in column (8) implies that, for home teams, a superstar – compared with a replacement-level player – improves his team's expected margin by +13 points (or, equivalently, 2.2 goals).

Taken together, the results in Tables 3 and 4 suggest that betting markets appear to appropriately price superstars' marginal productivity in the aggregate. According to the estimates we report in Table 3 column (4), over the course of a full season, a superstar will deliver his team an additional 1.9 wins (0.17*11 away games). According to the estimates we report in Table 4 column (4), if we allow that punters pool these superstar effects across home and away games, betting markets appear to believe that superstars deliver their teams an additional 1.9 wins (0.11*11 home games + 0.06*11 away games).

(iv) What Should a Superstar be Paid? Back-of-the-Envelope Calculations

As per Table 5, in 2018 (the final season in our dataset), superstars – players at the 95th percentile of the pay distribution – earned at least A $800,000 per season.[33] At the other end of the distribution, we define marginal or replacement-level players as individuals at the 45th percentile. This is based on several assumptions and stylised facts.

First, we assume that the distribution of pay across players within teams is similar to the pay distribution across all players in the league as a whole (i.e., each team has approx. two stars, who are paid a similar amount across teams, and each team has a large number of marginal players, who are paid a similar amount across teams). Second, we assume that performance or ability is strongly correlated with pay, so that teams effectively fill their line-ups in descending order of pay (i.e., teams field line-ups consisting of their 22 highest-paid, available players; as such, if a starting player is injured, he will be replaced by the highest-paid, non-starting player). As the average team's list has 703/18 ≈ 40 players, each team has 40–22 = 18 players that, absent injuries to starting players, do not make the line-up each week. As 18/40 = 0.45, we arrive at our definition of a marginal/replacement-level athlete as a player at 45th percentile of the pay distribution.

TABLE 5
Distribution of Player Pay in the Australian Football League (AFL)

Salary	No. of players	
A$0–60,000	1	
A$60,001–100,000	24	
A$100,001–200,000	159	
A$200,001–300,000	136	45th percentile
A$300,001–400,000	145	
A$400,001–500,000	88	
A$500,001–600,000	49	
A$600,001–700,000	46	
A$700,001–800,000	24	
A$800,001–900,000	15	95th percentile
A$900,001–1,000,000	10	
A$1,000,001–1,100,000	2	
A$1,100,001–1,200,000	1	
≥ A$1,200,001	3	
Total	703	

Notes: The distribution of AFL player pay for the 2018 season is shown. The salary cap during this season was A$12,450,000, and the average list size per team was approximately 40 players.
Source: AFL, 2018 Annual Report.

According to Table 5, teams typically pay marginal or replacement-level players somewhere between A$100,000 and A$300,000 per season. If we assume that teams overwhelmingly replace injured star players with the highest-paid (available) non-starters, we can conclude that, on average, teams in the AFL pay marginal players approximately A$250,000 per season. Drawing on these assumptions and stylised facts, we can use our estimates of a star player's marginal productivity to examine whether AFL teams appear to under or overpay for superstars.

If we take the average value of the coefficient on Away Star Injuries from Table 3 columns (1) to (4), we estimate that the loss of a superstar to injury lowers a team's likelihood of winning an upcoming away game by 15 percentage points. In a regular home and away season, a team plays 11 away games. This implies that compared with fielding a marginal/replacement-level player in each away game, by playing a superstar, a team,

[33] While not explicitly identified in Table 5, Lance Franklin and Jeremy Cameron reportedly each received salaries of approximately A$1,500,000 in 2018, making the two the highest-paid players in the AFL. Several media outlets provide yearly lists that purport to identify the AFL's highest-paid players, for example, Herald Sun, AFL.com. As the league does not disclose individual players' salaries, these outlets must rely on unofficial pay data provided by 'unnamed sources'. For the purposes of this study, we considered using data from these highest-paid player lists to identify superstar players in the AFL. However, concerns over the reliability of these data, our inability to verify the accuracy of the pay information and inconsistencies in data reporting (i.e., in some years, these lists identify the league's 30 highest-paid players, in other years these lists only identify the league's 12 highest-paid players) lead us instead to use a more transparent, objective, and data-driven approach for identifying superstars in our setting (i.e., the distribution of Champion Data's player ranking points). Putting these data issues to one side, we note that we nonetheless only find minor differences between lists of the league's top performing players – as per player performance ratings – and the league's highest-paid players – as per online sources. We argue that the overlap between these two sets (high-performing players and highly paid players) supports the validity of our approach for identifying superstar players in the AFL and speaks to the overall robustness of our paper's results.

FIGURE 1
Distribution of treatment effects from regression of game outcome on placebo injuries to star players on the away team.

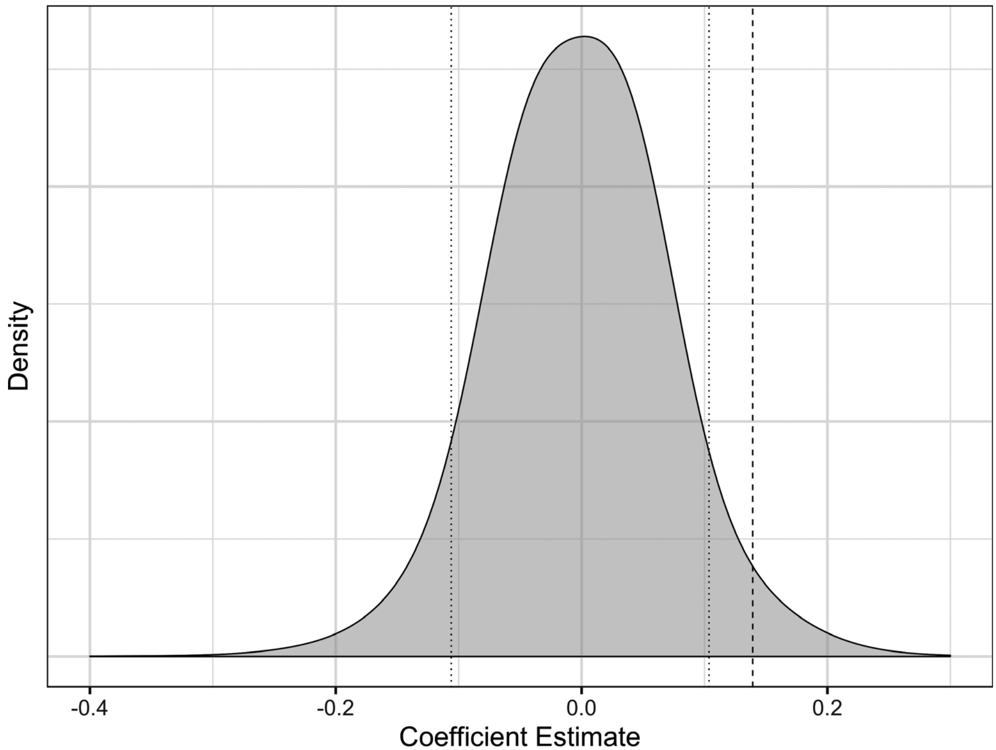

Note: The kernel density plot of the estimated treatment effects from the simulated regressions of game outcome on placebo injuries to star players on the away team (1000 replications) is shown. The average placebo effect is non-distinguishable from zero at the 5 per cent level (the dotted lines indicate the 95 per cent confidence interval). For comparison, the dashed line shows the estimate from our main regression using the actual number of star injuries to the away team (i.e., the estimated coefficient reported in Table 3, column 3).

on average, will win an additional 0.15*11 = 1.65 games per season.

Turning to player payments, teams pay the maximum amount permissible under the salary cap. Therefore, for the 2018 season teams aggregate spending on player payments was A $12,450,000. By construction, the average team in the AFL wins 50 per cent of games. This equates to 22/2 = 11 wins per season. This implies that the average team spends A$12,450,000/11 = A $1,132,000 per win on player payments.

Using these back-of-the-envelope calculations, we arrive at the following: if teams compensate superstars according to their productivity, stars should be paid the salary of a marginal/replacement player plus the total value of the star's marginal productivity. As we assumed earlier that the marginal AFL player earns A$250,000 per season, teams should pay superstar players A $250,000 + (1,132,000*1.65) = A$2,118,000 per season. On face value, our calculations suggest that superstar players are underpaid in the AFL.[34] As per Table 5, most star players are paid somewhere between A$800,000 and A

[34] As discussed in Section III.ii, as we only consider players' on-field productivity, our empirical strategy identifies lower-bound estimates of superstars' total productivity. As such, we argue that in practice superstars are likely to be even more heavily underpaid than our back-of-the-envelope calculations suggest.

$1,000,000, and no player is paid more than A $1,500,000.[35] This finding suggests that even the competition's highest-paid players are underpaid by approximately 30 per cent. Drawing on key features of the contracting environment in the AFL, we propose a number of potential explanations for this finding.[36]

As information in general is costly to collect and process and managers are subject to a range of well-documented cognitive limitations, teams may simply underestimate the productivity of star players relative to the productivity of marginal AFL players. AFL pundits and analysts make this argument somewhat regularly in the popular press (Fox Sports, 2020a). Research in economics shows that measuring and rewarding individual performance is difficult and costly and that firms in a variety of industries limit the adjustment of wage rates to relative productivity (Bishop, 1987; Bol, 2011). If we interpret our findings in light of this body of research, our results suggest that teams overestimate the relative productivity of marginal players and underestimate the relative productivity of star players, in turn, overpaying the former and underpaying the latter.

To also give some sense of the difficulty in identifying the marginal productivity of individual workers in team settings, we emphasise that our back-of-the-envelope pay calculations are based on somewhat noisy point estimates of superstars' marginal productivity. As shown in Table 3, while our estimates are statistically significant at conventional levels, our standard errors imply that our estimates nonetheless display a considerable degree of uncertainty. For instance, using the standard errors reported in column (2) to construct a 95 per cent confidence interval, our estimate of the marginal productivity of a superstar lies between 6.5 percentage points and 23.7 percentage points. Plugged back into our back-of-the-envelope calculations, these bounds imply that teams should pay superstar players somewhere between A$250,000 + (A$1,132,000*0.715) = A$1,059,830 per season and A$250,000 + (A$1,132,000*2.61) = A$3,204,520 per season.[37] This suggest two things: it is difficult to precisely estimate the optimal level of pay for superstars in the AFL (the upper bound calculation is approximately three times the lower bound calculation); but even so, our lower bound estimate of superstars' marginal value is well above the observed level of pay for the typical superstar in the AFL (albeit somewhat below the level of pay for the league's very highest paid players). However,

[35] To give some sense of how our assumptions influence these back-of-the-envelope calculations, we perform an untabulated sensitivity analysis that employs alternative assumptions. For instance, if we assume that replacement level players come from the 5th percentile (an assumption that is highly unlikely in practice but provides a useful lower bound), our calculations suggest that teams should pay superstars approximately A $100,000 + (A$1,132,000*1.65) = A$1,968,000. We also examine how our calculations change if we use the payroll per win for only the league's most successful teams (under this approach, we assume that the average team uses its payroll inefficiently, i.e., pays too much for a win). As the team at the top of the ladder in a given season usually wins approximately 17 games, we obtain that the best teams spend A$12,500,000/17 = A $735,000 per win. Using this figure for our back-of-the-envelope calculations, we find that teams should pay superstars A$250,000 + (A$735,000*1.65) = A $1,463,000. While we do observe that two players in the AFL earn this amount per season, our analysis suggests at a minimum that each team should be paying at least one player this amount.

[36] Another alternative explanation that we only discuss (and refute) here in passing is that teams pay superstars according to their marginal productivity but also anticipate that injuries or suspension will cause superstars to miss a certain number of games each season, that is, teams expect superstars to produce excess wins but a not a full season's 'worth'. So, how many games do teams need to expect superstars to miss for their current level of pay to be efficient? If we take it that a typical superstar earns A $900,000, we obtain that A$250,000 + (A $1,132,000*$x$) = A$900,000. If we solve for x, we get that $x = 0.5742$. As x is the number of excess wins a superstar generates in a season, we plug this back into our back-of-the-envelope calculation that uses the estimates from our empirical model, where y is the number of away games that a team anticipates a superstar will play each season: $0.15*y = 0.5742$. If we solve for y, we obtain $y = 3.8$. Rounding up, this suggests that for teams to pay superstars according to their marginal productivity, teams must anticipate that stars will only play approximately four away games each season (i.e., superstars will miss 11–4 = 7 away games each season due to injury or suspension). As data from the 2018 AFL Injury Report suggest that the average player misses 3.8 games per season due to injury, we argue that teams would need to expect superstars to miss an unreasonably large number of games for observed superstar pay to be efficient.

[37] To dissuade the reader of the notion that our upper-bound estimates are extreme or nonsensical, we note here that a number of high-profile media figures and player agents in the AFL argue that the competition's best players should be paid somewhere between A$3,000,000 and A$4,000,000 per season (Fox Sports, 2020b).

FIGURE 2
Distribution of treatment effects from regression of game margin on placebo injuries to star players on the away team.

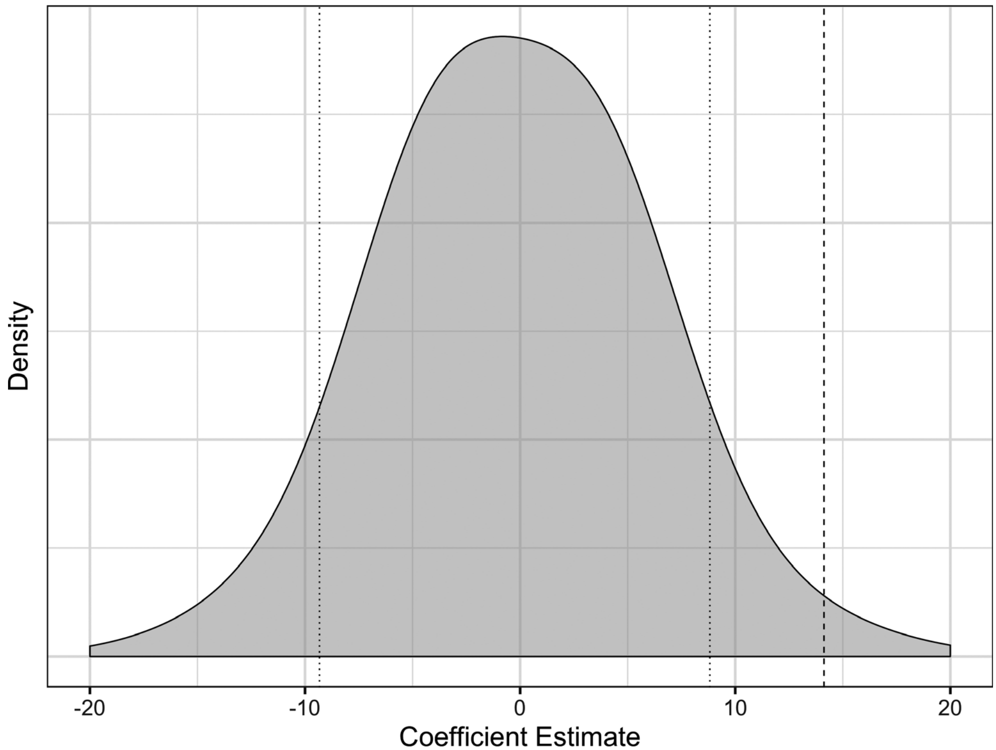

Note: The kernel density plot of the estimated treatment effects from the simulated regressions of game margin on placebo injuries to star players on the away team (1000 replications) is shown. The average placebo effect is non-distinguishable from zero at the 5 per cent level (the dotted lines indicate the 95 per cent confidence interval). For comparison, the dashed line shows the estimate from our main regression using the actual number of star injuries to the away team (i.e., the estimated coefficient reported in Table 3, column 7).

our upper bound estimate of superstars' marginal value is much higher than the observed level of pay for even the highest paid players in the AFL. We take this as further evidence that superstars in the AFL are likely underpaid in any given season (and at the very least highly unlikely to be overpaid in any given season).

Returning to potential explanations for why teams in the AFL may underpay superstars, we next note that teams may have a good sense of the relative productivity of star players but nonetheless underpay these players so as to limit harmful status effects and reduce disharmony and jealousy among teammates. In recent years, the popular press has attributed internal unrest at a number of teams to players' concerns over pay inequality and preferential treatment (ABC, 2020). Research in economics shows that in both sports and non-sports settings large relative pay disparities within organisations can also harm performance, lower job satisfaction and hamper collaboration among co-workers (Card et al., 2012). For example, Breunig et al. (2014) and Jewell (2017) find that salary inequality has a significantly negative effect on team success in US Major League baseball. Breza et al. (2018) uses a field experiment on Indian manufacturing workers to show that when co-workers' productivity is difficult to observe, pay inequality reduces output, lowers cooperation, and increases absenteeism. In our

TABLE 6
What Constitutes a Superstar? Regressions of Team Performance on Injuries to Players at the 90th Percentile

	Dependent variable							
	Win				Margin			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Home star injuries (90th)	−0.086	−0.101*	−0.078		−0.108*		−7.026	−8.079
	(0.067)	(0.056)	(0.049)		(0.059)		(5.534)	(5.321)
−6.062	−6.978							
(4.451)	(5.415)							
Away star injuries (90th)	0.071*	0.090**	0.099***		0.084*		7.880*	
	(0.037)	(0.037)	(0.035)		(0.047)		(4.343)	
9.836**			10.421**		9.225*			
(4.306)			(4.237)		(4.746)			
Home standing			−0.042***				−3.883***	
			(0.002)				(0.253)	
Away standing			0.036***				3.646***	
			(0.002)				(0.162)	
Home team FEs	No	Yes	Yes	No	No	Yes	Yes	No
Season FEs	No	Yes	Yes	No	No	Yes	Yes	No
Home team-season FEs	No	No	No	Yes	No	No	No	Yes
Observations	1133	1133	1124	1133	1133	1133	1124	1133
R^2	0.003	0.099	0.404		0.233		0.004	0.121
0.468	0.306							

Notes: *P < 0.1; **P < 0.05; ***P < 0.01. Estimates from ordinary least squares (OLS) regressions of game outcomes on our alternative superstar injury measures are presented. Table 1 defines each variable examined in this paper. Standard errors, reported in parentheses below the coefficient estimates, are clustered at either the home-team or home team-season level.

setting, if concerns over pay disparity and disharmony are sufficiently large, it may be efficient for teams to overpay marginal players and underpay stars.[38]

'Social outrage' and subsequent labour market regulations may also 'force' teams to underpay superstars relative to their marginal productivity. Research suggests that in the context of CEO pay persistent outrage expressed by politicians, the press and media, labour unions, and the general public may act to constrain executive compensation (Murphy & Jensen, 2018). The AFL requires teams to pay players a minimum salary that is tied to the players' level of professional experience. Under the hard salary cap enforced by the AFL, this minimum wage may work to compress star pay and force teams to transfer rents to marginal players. That is, teams may wish to pay marginal players less and pay star players more but cannot do so because of binding ceilings and floors on player pay (Booth et al., 2012).[39] Through the league's collective bargaining agreement, union influence, in the

[38] By censoring the right-tail of the pay distribution, the AFL's salary cap likely amplifies status effects and jealousy that stem from inequality in relative pay among teammates (i.e., the salary cap constrains pay and so players compete for status on the basis of their rank in the pay distribution). These effects are also likely to be especially pronounced in our setting given that highly competitive, 'type A' personalities typically select into careers in professional sports (Hopkins & Kornienko, 2009; Tran & Zeckhauser, 2012) and that research shows that peer-performance comparisons strongly motivate professional athletes (Eyring et al., 2021).

[39] Clubs' use of third-party deals to pay players outside the salary cap is consistent with the notion that, absent the salary cap, teams would pay superstar players more than they currently do. As an example of such a practice, Carlton notoriously arranged for a club sponsor to pay star midfielder Chris Judd A $250,000 per season for his services as a 'corporate ambassador'. While the AFL has now banned third-party deals, our findings suggest that, in the presence of a salary cap, such arrangements may improve efficiency in the sense that they allow superstars to be paid their marginal product.

TABLE 7

What Constitutes a Superstar? Regressions of Betting Market Expectations on Injuries to Players at the 90th–95th Percentiles

	Dependent variable							
	Win				Margin			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Home star injuries (90th–95th)	−0.051	−0.078	−0.079	−0.091	−4.271	−6.996	−7.323	−7.475
	(0.065)	(0.064)	(0.052)	(0.078)	(6.813)	(7.558)	(5.923)	(7.585)
Away star injuries (90th–95th)	0.044	0.066	0.084*	0.048	6.005	7.699	9.020	6.341
	(0.047)	(0.045)	(0.048)	(0.059)	(5.075)	(5.266)	(6.063)	(5.878)
Home standing			−0.042***				−3.878***	
			(0.002)				(0.253)	
Away standing			0.036***				3.656***	
			(0.002)				(0.163)	
Home team FEs	No	Yes	Yes	No	No	Yes	Yes	No
Season FEs	No	Yes	Yes	No	No	Yes	Yes	No
Home team-season FEs	No	No	No	Yes	No	No	No	Yes
Observations	1133	1133	1124	1133	1133	1133	1124	1133
R^2	0.001	0.096	0.402	0.230	0.001	0.118	0.466	0.304

Notes: $^{*}P < 0.1$; $^{**}P < 0.05$; $^{***}P < 0.01$. Estimates from ordinary least squares (OLS) regressions of game outcomes on our alternative superstar injury measures are presented. Table 1 defines each variable examined in this paper. Standard errors, reported in parentheses below the coefficient estimates, are clustered at either the home-team or home team-season level.

form of the AFL Players' Association, may also play a role in this process (ABC, 2017).

Another explanation for our findings is that teams use contracts that backload compensation and smooth players' earnings over the course of their career. In recent years, AFL teams have signed a number of high-profile players to long-term contracts (e.g., 7–10-year commitments). Under these contracts, star players are typically paid either a salary of upwards of A$1,000,000 per season that is fixed for the duration of the contract or a salary that progressively increases season on season for the duration of the contract (The Age, 2013; AFL, 2017). As such, teams may employ long-term contracts that underpay stars in their prime and overpay stars when they are past their prime. As per predictions from contract theory, if the net present value of such a long-term contract is the same as the net present value of a contract that pays a star player a yearly salary based on their expected productivity in that season (i.e., stars get paid less and less as their productivity declines over time), superstars in the AFL may not in fact be underpaid when total earnings across their entire contract – rather than in a given season – are taken into consideration (Kotlikoff & Gokhale, 1992). Research in economics on agency models shows that firms in a range of industries may use these types of long-term, backloaded contracts to smooth employee earnings and mitigate shirking (Becker & Stigler, 1974; Lazear, 1981). In pro sports, teams may be drawn to these long-term contracts as they reduce pay inequalities across players and provide flexibility for dealing with hard salary caps. Pre-empting rising pay levels (e.g., due to expected increases in the salary cap), teams may also use long-term contracts to lock in players on nominally high fixed salaries and then rely over time on 'inflationary' forces to lower player payments in 'real' terms.[40]

A final explanation for our findings is that players may derive utility from monetary and non-monetary factors and that superstars accept on-field success, that is, winning premierships, as

[40] As an example of such a contract, the Brisbane Lions famously signed Alastair Lynch to a 10-year contract in 1994. While in the early years of this deal Lynch was among the league's highest paid players, the rapid commercialisation of the AFL and associated increases in the salary cap during the late 1990s meant that later in his career Lynch was very likely underpaid relative to his peers.

a form of 'payment in kind'. If superstars derive relatively more utility from team success than non-superstars, then superstars might be willing to trade off their own salary to increase their team's chances of winning a premiership, that is, the superstar takes a pay cut so that his club can use the cap space to re-sign his teammates or recruit players from successful rival teams. Anecdotal evidence from media coverage at the time suggests that in the early 2000s star players at the Brisbane Lions took 'pay cuts' so that the team could enjoy continued on-field success (The Age, 2003). Such trade-offs also need not be entirely voluntary in that superstars may face pressure from their peers to take a pay cut for the 'good of the team'.[41] While these non-monetary factors may explain why successful teams can underpay superstars in equilibrium, they are less useful for explaining why superstars at consistently unsuccessful teams are underpaid.

V Robustness Checks

We conduct a number of robustness checks that evaluate the plausibility of our empirical design. To identify superstars' marginal productivity, injury-induced line-up changes between games must be plausibly exogenous. To support this assumption, we discussed in Section IV.ii how Table 2 shows that superstar injuries appear to be unrelated to a range of observable pretreatment outcomes. We now build on this evidence by performing a set of placebo tests.

First, we conduct placebo analyses whereby for each game in our sample we randomly assign injury-induced line-up changes to teams' star players. To do so, we sample without replacement from the empirical distribution of the number of injury-induced line-up changes observed in our sample. Using these placebo injuries, we then re-estimate Equations (1) and (2), that is, our regressions that estimate superstar players' marginal productivity. Using this procedure, we estimate 1000 simulated regressions that we then use to generate empirical distributions of these placebo treatment effects. If we observe that the average placebo effect is statistically significant, injury-induced line-up changes are unlikely to be exogenous and our main results are likely to be biased.

In Figures 1 and 2, we plot these distributions and show that the mean placebo treatment effects are non-distinguishable from zero at the 5 per cent level. For comparison, we also plot the estimate from our main regression results that use the actual number of star injuries to the away team (i.e., the estimated coefficients reported in Table 3, column 3, and Table 3, column 7). We present the plots reported in Figures 1 and 2 as evidence in support of our identification strategy.

Second, we perform an alternative placebo test where we weaken our definition of superstar players and test whether we continue to observe superstar effects for these 'good-but-not-great' players. We motivate these tests using the logic that our main results are more likely to be spurious if we also observe superstar effects for these 'non-superstars'.

In Table 6, we report results from estimations of Equations (1) and (2) where we define a superstar as a player in the 90th percentile of the distribution of Champion Data's player ranking points. We observe that when using this more generous definition of a superstar, our results decrease in magnitude and statistical significance. Specifically, the coefficient on away star injuries is approximately half the size of the estimate reported in Table 3. As the estimation procedure used to generate the results reported in Table 6 pools together superstars and 'good-but-not-great' players, this finding is consistent with non-superstars 'diluting' estimates of superstars' abnormal marginal productivity.

We shed further light on this finding by examining what happens to our estimates when we define a superstar as a player in the 90th–95th percentiles of the performance distribution (i.e., when we estimate the marginal productivity of only the 'good-but-not-great' players). In Table 7, we fail to observe that teams are statistically significantly less likely to lose an upcoming game if these players withdraw due to injury from the teams' line-up. In short, we do not observe estimates consistent with superstar effects for players that 'just' fail to meet the definition of a superstar that we use in our main analyses, that is, superstar effects rapidly vanish as we move inward from the right tail of the performance distribution.

In summary, Tables 6 and 7 provide evidence that our superstar effects only show up for

[41] For example, Joel Selwood reportedly agreed to a large pay cut so that the Geelong Cats could afford to sign superstar forward Jeremy Cameron from the Greater Western Sydney Giants. Similarly, Carlton's Patrick Cripps recently accepted a reduced salary to provide his team with increased cap space and, in turn, a better chance at premiership success in the future.

workers at the very end of the right tail of the performance distribution – where the literature suggests we should expect to detect them. As such, these robustness checks both validate our measurement approach and support our identification strategy, that is, if injury-induced line-up changes are endogenous and our empirical design is picking up the effect of some unobserved confounder, than this same confounder should also produce spurious results when we use these alternative superstar definitions. We also note that Tables 6 and 7 provide insight into the question of 'who' should be considered a superstar in the context of the AFL. As we show, superstars in the AFL are rare, and the performance gap between very good and great players is large.[42]

In this section, we also evaluate the construct validity our superstar measure. While many different metrics can be used to identify superstars in the AFL, each approach has its own advantages and disadvantages. To validate our method, we document how players in the 95th percentile of the ratings distribution perform on four alternative superstar indicators. While we do not expect our superstar measure to perfectly comove with these alternative proxies, we do expect considerable overlap as each broadly captures some dimension of the same construct, that is, superstar performance.

First, we examine how our superstar players perform in their teams' Best and Fairest voting. While voting systems vary, all teams in the AFL award a Best and Fairest award each season. As such, Best and Fairest voting is a highly representative method for identifying superstar players (i.e., it requires that superstars are drawn from all teams, even those that perform very poorly). Second, we examine how our superstar players perform in the Brownlow Medal, the AFL's top individual honour. The Brownlow Medal is determined by votes cast by the officiating umpires after each game. Critics of the Brownlow Medal argue that the award favours midfielders and fails to acknowledge high-performing players in other positions. Third, we examine how frequently our superstars appear in the All-Australian team. Each season, the AFL appoints a panel of experts to select an All-Australian team. Unlike the Brownlow Medal, All-Australian honours are distributed across the full set of playing positions on an AFL team. Finally, we examine how our superstars perform in the AFL Players Association's Most Valuable Player (MVP) award. The MVP is a 'peer award' that is determined by players' votes. Unlike the Brownlow Medal, individuals remain eligible for the MVP even if suspended and players from a wide range of playing positions have won the award.

Across the period 2016–18, we do indeed find that our superstar measure is strongly associated with these alternative superstar indicators. First, we find that, on average, players in the 95th percentile of the ratings distribution account for approximately two-thirds of all top-two places in teams' Best and Fairest voting each season. We find that our superstars are less likely to be Best and Fairest winners from teams that performed very poorly in a given season (i.e., players over whom there is the most debate over whether they qualify as 'superstars'). Second, we find that a player in the 95th percentile won each of the Brownlow Medals awarded during our sample period. We also find that, on average, our superstars account for nine of the top-10 places in Brownlow Medal voting each season. Third, we find that, on average, players in the 95th percentile account for approximately 12 of the 22 spots in the All-Australian team that is selected each season. We observe that our superstars dominate selection for All-Australian midfield, half-back and half-forward positions. We also find for a majority of seasons in our sample that a player in the 95th percentile of the ratings distribution received the added distinction of being made the sole 'captain' of the All-Australian team. Finally, we find that our superstars account for all top two places in the AFL MVP award. In summary, these descriptive findings support the overall validity of our approach for identifying superstar players in the AFL and strongly suggest that our main results are unlikely to be driven by idiosyncrasies of our specific superstar measure.

VI Conclusion

In this paper, we use game-level data from the AFL to examine superstar workers' productivity and pay. We exploit injury-induced changes to AFL team line-ups between games to produce robust estimates of superstar players' marginal productivity.

[42] As mentioned in the introduction, while it would be insightful to estimate the marginal productivity of players at the 99th percentile, we simply do not observe enough injuries among this very small group of elite players to implement our empirical strategy.

We find that compared with replacement-level players, superstars in the AFL increase their teams' likelihood of winning away games by approximately 15 percentage points. While we show that betting markets appear to appropriately price superstars' marginal productivity in the aggregate, we present back-of-the-envelope calculations that suggest that AFL teams underpay superstar players by at least 30 per cent. While payroll data indicate that the AFL's highest-paid players are paid approximately A$1,200,000–1,500,000 per season, our estimates suggest that the league's very best players should be paid in excess of A$2,000,000 per season. We discuss how this result may be explained by inaccuracies in managers' performance evaluations, distortions from labour market regulations, the use of long-term back-loaded contracts, teams' attempts to mitigate the harmful effects of pay disparities among co-workers, and on-field success as a form of payment-in-kind.

Our study extends the empirical economics literature on superstars by using a research design that exploits plausibly exogenous variation in firms' workforce composition to generate unbiased estimates of superstars' marginal productivity. Our study also contributes to research on the economic behaviour of professional sports organisations by linking unbiased estimates of superstar players' marginal productivity to the distribution of player pay in the AFL.

Our paper also has implications for future research that uses data from professional sports to examine broader questions in labour and personnel economics. For instance, our research design is well suited to studying superstar productivity and pay in other professional sports settings where injuries act as plausibly exogenous shocks to team composition, for example, rugby league or union, basketball, soccer. Given that the relatively small number of superstar injuries in our dataset means that we are unable to examine, with any reasonable level of statistical precision, how superstar effects vary in the cross-section, future research could also collect data for additional AFL seasons and extend our study by using this larger sample to explore how superstar productivity and pay varies by superstars' playing position or by teams' talent distribution (e.g., Is a forward superstar more important than a star ruckman? Are superstar teammates complements or substitutes?). More generally, researchers can also use our injury-based empirical design to study how exogenous changes in team composition affect various organisational outcomes of interest to labour economists (e.g., accumulation of team-specific human capital, disruptions to on-the-job training, off-field financial performance, contract renegotiations, etc).

Acknowledgements

Open access publishing facilitated by The University of Melbourne, as part of the Wiley – The University of Melbourne agreement via the Council of Australian University Librarians.

Conflict of interest

The authors have declared no conflicts of interest.

Supporting Information

Additional Supporting Information may be found in the online version of this article (https://onlinelibrary.wiley.com/doi/10.1111/1475-4932.12665):

Data S1: Codebook, data file, and R script for descriptive statistics, model estimation, and figures

REFERENCES

ABC. (2017), *AFL Pay Rise*. Accessed June 2021. Available from: https://www.abc.net.au/news/2017-06-23/afl-pay-rise-richard-hinds-column/8642772.

ABC. (2020), *New Bulldog Treloar Reveals Hurt From Collingwood Snub*. Accessed June 2021. Available from: https://www.abc.net.au/news/2020-11-13/new-bulldog-treloar-reveals-hurt-from-collingwood-snub/12881510.

AFL. (2017), *Dusty Mega Deal Done: Superstar a Tiger for Life*. Accessed June 2021. Available from: https://www.afl.com.au/news/145426/dusty-mega-deal-done-superstar-a-tiger-for-life.

AFL. (2018), *2018 Annual Report*. Available from: https://resources.afl.com.au/afl/document/2019/12/05/aad564ee-58ac-4f16-97de-8d930a5237ef/2018_Digital_AR_Book-min.pdf.

AFL. (2019), *2019 Annual Report*. Accessed June 2021. Available from: https://resources.afl.com.au/afl/document/2020/03/18/925fd047-a9b6-4f7d-8046-138a56ba36f4/2019-AFL-Annual-Report.pdf.

Angrist, J.D. and Pischke, J.-S. (2008), *Mostly Harmless Econometrics*. Princeton University Press, Princeton.

Azoulay, P., Graff Zivin, J.S. and Wang, J. (2010), 'Superstar Extinction', *The Quarterly Journal of Economics*, **125**, 549–89.

Becker, G.S. and Stigler, G.J. (1974), 'Law Enforcement, Malfeasance, and Compensation of Enforcers', *The Journal of Legal Studies*, **3**, 1–18.

Berri, D.J., Schmidt, M.B. and Brook, S.L. (2004), 'Stars At the Gate: The Impact of Star Power on NBA Gate Revenues', *Journal of Sports Economics*, **5**, 33–50.

Bishop, J. (1987), 'The Recognition and Reward of Employee Performance', *Journal of Labor Economics*, **5**(4 Part 2), S36–56.

Blackham, J. and Chapman, B. (2004), 'The Value of Don Bradman: Additional Revenue in Australian Ashes Tests', *Economic Papers: A Journal of Applied Economics and Policy*, **23**, 369–85.

Blainey, G. (2010), *A Game of Our Own: The Origins of Australian Football*. Black Inc., Melbourne.

Bol, J.C. (2011), 'The Determinants and Performance Effects of Managers' Performance Evaluation Biases', *The Accounting Review*, **86**, 1549–51.

Booth, R., Brooks, R. and Diamond, N. (2012), 'Player Salaries and Revenues in the Australian Football League 2001–2009: Theory and Evidence', *The Economic and Labour Relations Review*, **23**, 39–54.

Borland, J. and Booth, R. (2014), *AFL Team Performance and Expenditure: Is It How Much Money Your Club Spends, or How They Spend It, That Matters?* Working paper, University of Melbourne, Melbourne. Available from https://docs.google.com/file/d/0B_H1wGTm98W3M2xqSGVHVlJKVnc/edit?resourcekey=0-h9hIApHriAtAN3iSJxszQw.

Breunig, R., Garrett-Rumba, B., Jardin, M. and Rocaboy, Y. (2014), 'Wage Dispersion and Team Performance: A Theoretical Model and Evidence from Baseball', *Applied Economics*, **46**, 271–81.

Breza, E., Kaur, S. and Shamdasani, Y. (2018), 'The Morale Effects of Pay Inequality', *The Quarterly Journal of Economics*, **133**, 611–63.

Card, D., Mas, A., Moretti, E. and Saez, E. (2012), 'Inequality At Work: The Effect of Peer Salaries on Job Satisfaction', *American Economic Review*, **102**, 2981–3003.

Clarke, S.R. and Norman, J.M. (1995), 'Home Ground Advantage of Individual Clubs in English Soccer', *Journal of the Royal Statistical Society: Series D (The Statistician)*, **44**, 509–21.

Dang, T.M., Booth, R., Brooks, R. and Schnytzer, A. (2015), 'Do TV Viewers Value Uncertainty of Outcome? Evidence from the Australian Football League', *Economic Record*, **91**, 523–35.

Depken, C.A. (2000), 'Wage Disparity and Team Productivity: Evidence from Major League Baseball', *Economics Letters*, **67**, 87–92.

Eyring, H., Ferguson, P.J. and Koppers, S. (2021), 'Less Information, More Comparison, and Better Performance: Evidence from a Field Experiment', *Journal of Accounting Research*, **59**, 657–711.

Fox Sports. (2020a), *Ordinary Players Paid Too Much*. Accessed June 2021. Available from: https://www.foxsports.com.au/afl/afl-2020-salary-cap-total-player-payments-ordinary-players-paid-too-much-jordan-lewis-leigh-montagna-fox-footy-podcast/news-story/df7bc15b771eee365f016ea1711dc29a.

Fox Sports. (2020b), *Top Players Deserve A$3 Million*. Accessed June 2021. Available from: https://www.foxsports.com.au/afl/afl-2020-top-players-deserve-3-million-a-year-middle-players-paid-three-times-too-much-afl-salary-cap-tpp-average-wage/news-story/874317ad366820cca7a4021ca5bce65a.

Gabaix, X. (2016), 'Power Laws in Economics: An Introduction', *Journal of Economic Perspectives*, **30**, 185–206.

Gabaix, X. and Landier, A. (2008), 'Why Has CEO Pay Increased So Much?', *The Quarterly Journal of Economics*, **123**, 49–100.

Greene, W. (2004), 'The Behaviour of the Maximum Likelihood Estimator of Limited Dependent Variable Models in the Presence of Fixed Effects', *The Econometrics Journal*, **7**, 98–119.

Groysberg, B., Lee, L.-E. and Nanda, A. (2008), 'Can They Take It With Them? The Portability of Star Knowledge Workers' Performance', *Management Science*, **54**, 1213–30.

Hausman, J.A. and Leonard, G.K. (1997), 'Superstars in the National Basketball Association: Economic Value and Policy', *Journal of Labor Economics*, **15**, 586–624.

Herald Sun. (2013), *Hawthorn Has Let Lance Franklin Become a Sydney Swan With The AFL Approving the Mega Deal*. Accessed June 2021. Available from: https://www.heraldsun.com.au/sport/afl/hawthorn-has-let-lance-franklin-become-a-sydney-swan-with-the-afl-approving-the-mega-deal/news-story/6718e7dddf4a6e2b60ce7cebf04859f1.

Hopkins, E. and Kornienko, T. (2009), 'Status, Affluence, and Inequality: Rank-based Comparisons in Games of Status', *Games and Economic Behavior*, **67**, 552–68.

Horrace, W.C. and Oaxaca, R.L. (2006), 'Results on the Bias and Inconsistency of Ordinary Least Squares for the Linear Probability Model', *Economics Letters*, **90**, 321–7.

Hrysomallis, C. (2013), 'Injury Incidence, Risk Factors and Prevention in Australian Rules Football', *Sports Medicine*, **43**, 339–54.

Humphreys, B.R. and Johnson, C. (2020), 'The Effect of Superstars on Game Attendance: Evidence from the NBA', *Journal of Sports Economics*, **21**, 152–75.

Jane, W.-J. (2016), 'The Effect of Star Quality on Attendance Demand: The Case of the National Basketball Association', *Journal of Sports Economics*, **17**, 396–417.

Jewell, R.T. (2017), 'The Effect of Marquee Players on Sports Demand: The Case of US Major League Soccer', *Journal of Sports Economics*, **18**, 239–52.

Kahn, L.M. (2007), 'Markets: Cartel Behavior and Amateurism in College Sports', *Journal of Economic Perspectives*, **21**, 209–26.

Kotlikoff, L.J. and Gokhale, J. (1992), 'Estimating A Firm's Age–Productivity Profile Using the Present Value of Workers' Earnings', *The Quarterly Journal of Economics*, **107**, 1215–42.

Lazear, E.P. (1981), ‘Agency, Earnings Profiles, Productivity, and Hours Restrictions’, *The American Economic Review*, **71**, 606–20.
Lucifora, C. and Simmons, R. (2003), ‘Superstar Effects in Sport: Evidence from Italian Soccer’, *Journal of Sports Economics*, **4**, 35–55.
Malmendier, U. and Tate, G. (2009), ‘Superstar CEOs’, *The Quarterly Journal of Economics*, **124**, 1593–638.
McCall, A., Carling, C., Nedelec, M., Davison, M., Le Gall, F., Berthoin, S. and Dupont, G. (2014), ‘Risk Factors, Testing and Preventative Strategies for Non-contact Injuries in Professional Football: Current Perceptions and Practices of 44 Teams From Various Premier Leagues’, *British Journal of Sports Medicine*, **48**, 1352–7.
Murphy, K.J. and Jensen, M.C. (2018), ‘The Politics of Pay: The Unintended Consequences of Regulating Executive Compensation’, *Journal of Law, Finance, and Accounting*, **3**, 189–242.
Pinnuck, M. and Potter, B. (2006), ‘Impact of On-field Football Success on the Off-field Financial Performance of AFL Football Clubs’, *Accounting & Finance*, **46**, 499–517.
Rosen, S. (1981), ‘The Economics of Superstars’, *The American Economic Review*, **71**, 845.
Rosen, S. and Sanderson, A. (2001), ‘Labour Markets in Professional Sports’, *The Economic Journal*, **111**, F47–68.
Scully, G.W. (1974), ‘Pay and Performance in Major League Baseball’, *The American Economic Review*, **64**, 9.
SEN. (2020), *Should the AFL Allow Clubs to Pay a Marquee Player Outside of the Salary Cap?*. Accessed June 2021. Available from: https://www.sen.com.au/news/2020/05/21/should-the-afl-allow-clubs-to-pay-a-marquee-player-outside-of-the-salary-cap/.
Stefani, R. and Clarke, S. (1992), ‘Predictions and Home Advantage for Australian Rules Football’, *Journal of Applied Statistics*, **19**, 251–61.
Sydney Swans. (2013), *The Papers Need to Get Over It’: Sydney Scrutiny Wearing Thin for Buddy*. Accessed June 2021. Available from: https://www.sydneyswans.com.au/news/717790/the-papers-need-to-get-over-it-sydney-scrutiny-wearing-thin-for-buddy.
The Age. (2003), *Lions to Be Hit Hard By Cap Cuts*. Accessed June 2021. Available from: https://www.theage.com.au/sport/afl/lions-to-be-hit-hard-by-cap-cuts-20030701-gdvywo.html.
The Age. (2013), *Revealed Buddy Frankling to Profit Heavily in Later Years of Megadeal*. Accessed June 2021. Available from: https://www.theage.com.au/sport/afl/revealed-buddy-franklin-to-profit-heavily-in-later-years-of-megadeal-20131009-2v8qk.html.
Tran, A. and Zeckhauser, R. (2012), ‘Rank as an Inherent Incentive: Evidence from a Field Experiment’, *Journal of Public Economics*, **96**, 645–50.
Wooldridge, J.M. (2010), *Econometric Analysis of Cross Section and Panel Data*. MIT Press, Cambridge.