# Peer Health Shocks and Labor Supply

Joshua C. Martin

Vanderbilt University

November 14, 2024

Click here for the most up-to-date version

**Abstract**

---

I provide novel evidence on how workers respond to peer health shocks within high-risk occupations by leveraging two nested natural experiments within professional hockey and American football. First, I compare differences in labor supply between characteristically similar athletes who differ only in their exposure to a colleague who died of chronic traumatic encephalopathy (CTE) – a deadly neurological disease causally linked to continued workplace participation. Though the information about these deaths is widely publicized, I find that their occurrence differentially increases the probability for former teammates to retire. This effect is greater for those with longer periods spent as teammates and diminishes with time since they were last on the same team. Second, I leverage quasi-random differences in the monetary compensation that workers would forgo upon retiring at the time of this peer health shock. I show these retirements are highly responsive to opportunity costs – estimating that teams would have to increase worker compensation $6 million to prevent their exit. Remaining treated workers display a heightened sensitivity to health risks by exchanging salary for larger signing bonuses and shorter contracts in their subsequent employment negotiations. The finding that labor supply decisions are highly responsive to the health status of peers suggests that workers substantially underestimate utility loss from work-related health damages even in environments where such risks are highly publicized.

---

**Keywords:** labor supply, health shocks, peer effects, occupational risk, compensating wage differentials, opportunity costs
**JEL Codes:** J28, J24, J33, I12, Z22

*"Concussions are part of the profession, an occupational risk... like a steelworker who goes up 100 stories, or a soldier."*

- Elliot J. Pellman, Chairman of the NFL's Mild Traumatic Brain Injury Committee

email: joshua.martin@vanderbilt.edu
I am thankful to Kitt Carpenter, Marcus Dillender, Brad Humphreys, Joshua Hall, Bryan McCannon, Brian Beach, Alexander Cardazzi, Hasan Shahid, Laura Nettuno, João Tampellini and Tyler Schumacher for their comments on earlier versions of this paper including seminar participants and discussants in the Economics Department at Vanderbilt University, the Economics Department at Lafayette College, the Sports Management faculty at University of Georgia, the 2024 American Society of Health Economics Meeting, and the 2024 Western Economic Association meeting. Special thanks are also due to two individuals without whom this project would not exist. First, this paper builds on work conducted jointly with my former student, Logan Tucker. His passion for the setting and careful articulation of the plausible economic forces driving labor supply through peer effects helped to pave the way for this project. Second, I owe much to Clay Collins, whose encouragement, aid in idea generation and rich knowledge of the setting were instrumental in the completion of this draft. His patience and generous time investments in helping me understand the complicated phenomenon of labor supply in a sports setting have deepened my understanding of Economics.

# 1 Introduction

The factors that influence how individuals perceive and respond to changes in risk have long been a central question for social scientists. While much research has examined the roles of institutions, macroeconomic forces, and the natural environment in shaping these preferences (Doepke & Zilibotti, 2008; Falk et al., 2018; Galor & Özak, 2016; Hanaoka et al., 2018; Malmendier & Nagel, 2011), a growing body of literature has begun to focus on the impact of more localized forces, particularly social networks (Dohmen et al., 2012). In this paper, I contribute to this literature by examining specific mechanisms through which peers shape individuals' risk perceptions and behaviors.

To address this research question, I examine a uniquely advantageous setting – professional sports. The intense popularity of sports, along with meticulous historical data-keeping practices, enables the construction of detailed peer networks by tracking athletes throughout their professional, collegiate, and even high school careers. The draft process in professional sports leagues also reduces concerns about selection endogeneity with respect to peers, a well-known challenge in empirical research on peer effects (Manski, 1993). Rich data on player characteristics such as age, skill specializations, and productivity further allow for the use of matching methods to construct plausible counterfactual groups. Finally, this unique setting provides clear identification of work-related hazards, overcoming issues of ambiguity that are likely present in other contexts.

I leverage two natural experiments within this sports setting. First, I exploit quasi-random variation in athletes' exposure to peers diagnosed with chronic traumatic encephalopathy (CTE), a fatal neurological disease causally linked to participation in violent, collision sports. Using a matched difference-in-differences design with staggered treatment timing, I compare characteristically similar athletes who differ only in their exposure to a former teammate who died from this work-related injury. This approach allows me to estimate the differential impact of these health shocks on former teammates' subsequent labor supply decisions based on the timing of the CTE deaths. Second, by utilizing detailed contract data, I analyze differences in the remaining unpaid compensation on these athletes' contracts at the time of these health shocks to estimate the monetary value of this disamenity through these workers' revealed preferences.

Examining the complete set of publicly confirmed CTE diagnoses among former athletes in the National Football League (NFL) and National Hockey League (NHL), the results indicate that a former teammate's CTE-related death significantly raises the probability of workers exiting the profession by 2.2 to 4.0 percentage points (8–23%). These effect sizes are substantial, accounting for approximately one-tenth to one-fourth of a cohort exiting due to these deaths. I identify two key channels through which this elevated exit rate arises. First, the probability of exit increases with personal familiarity, measured by the duration spent as teammates. Second, these effects diminish as the time since individuals were last colleagues increases. Models incorporating both mechanisms (athletes who were longtime teammates and recently experienced a peer's CTE-related death) reveal that their interaction amplifies the impact, producing a combined effect greater than the sum of their individual contributions. These findings suggest that the immediacy and personal relevance of informational health shocks are pivotal in shaping workers' perceptions and responses to occupational health risks. Additionally, the results help explain variations in labor supply responses across sports; hockey players, in particular, were more likely to exit following a peer health shock, likely due to smaller team sizes that foster closer familiarity, historically coincidental shorter time intervals between CTE deaths, and significantly weaker labor demand, which limits workers' bargaining power.

To rule out the possibility that these findings are driven by factors unrelated to changes in risk perception, I substitute the indicator for exposure to a former teammate who died of CTE with variables capturing exposure to a former teammate who died of non-workplace-related causes. For this analysis, I focus on the two most common causes of death among prime-age workers: car accidents and diseases such as cancer. Estimates from these models reveal precise null impacts on labor supply, suggesting that factors such as grief are not driving the observed results.

I then investigate whether wages for treated workers increased to offset the impact of the peer health shock on exit rates. Although I find economically meaningful positive effects of treatment on total monetary compensation, these effects are not statistically significant, likely because workers at the margin of retirement near the time of the shock have limited bargaining power in contract negotiations. However, this overall effect masks substantial heterogeneity in the structure of subsequent contracts. Treated workers exchange reductions in salary for increased guaranteed compensation and shorter contract durations. This pattern among remaining workers is consistent with heightened sensitivity to health risks, suggesting that treated workers were willing to stay only if they could secure stronger financial guarantees for their future income.

Finally, I propose an alternative method for estimating compensating wage differentials based on opportunity costs. This approach leverages a unique identification strategy based on how contracts are restructured to meet team-level salary cap restrictions: individuals with similar total compensation are paid different annual amounts. As a result, athletes with comparable overall compensation and productivity may face vastly different opportunity costs at the time of exposure. This strategy resembles the approach by Imbens et al. (2001), where "among lottery winners, the magnitude of the prize is randomly assigned". By distinguishing between salary and guaranteed income—representing variable and fixed costs in retirement decisions—the contract data allows for precise calculation of the labor income athletes would forgo upon retirement in any given period. My findings show that retirement decisions are highly sensitive to these opportunity costs. Estimates from a difference-in-differences model, comparing individuals with similar annual wages but differing in forgone income at the time of treatment, suggest that former teammates of athletes who died of CTE would require $6 million to offset this disamenity and $1 million to be indifferent between exiting and staying in the profession.

This paper demonstrates that peer health shocks have a strong impact on workers' labor supply, a finding that is surprising for several reasons. First, professional athletes earn exceptionally high wages, and firm entry into these leagues is highly restricted, resulting in a large excess supply of individuals aspiring to join. This intense competition has led to early and substantial investments in sport-specific skills, which are not easily transferable to other markets. Consequently, workers in these professions have limited leverage to credibly threaten to leave voluntarily. Second, the persistence of these labor supply responses highlights the role of peer effects. Despite widespread media coverage following any diagnosis of CTE in athletes or repeated concussions among superstars, which also reaches individuals in the control group, the observed effects are still pronounced. This suggests that the true effects of peer health shocks on labor supply is even stronger than the findings presented in this paper indicate. This finding aligns with the notion that labor supply responses in this context are driven by the salience of the peer health shock, rather than an information-based updating process (Bleemer & Zafar, 2018; Dessaint & Matray, 2017).

Taken together, these results provide strong evidence of two key phenomena. First, individuals are significantly more responsive to informational signals, such as health shocks, when they have a personal connection to the messenger or affected individual, with recency and familiarity emerging as key determinants. Second, these findings suggest that workers significantly underestimate the utility loss associated with work-related health damages, even in environments where such risks are highly publicized and widely discussed. This underestimation can lead to suboptimal career decisions and increased exposure to workplace risks that are insufficiently compensated, ultimately endangering individual health and well-being. The broader implications extend to long-term aggregate productivity, as these misjudgments can reduce workforce stability and efficiency over time. These results highlight the importance of more effective communication about occupational health risks, particularly through credible messengers who have personal connections to workers.

## 2 Literature Review

Social scientists have long studied the effects of peer spillovers on individuals' actions and attitudes. A major force contributing to the sheer volume of this literature was the innovation in empirical methodology for measuring these phenomena introduced by Manski (1993). More recently, the proliferation of methodologies from the causal inference revolution, coupled with a more widespread availability of panel data, has led researchers to increasingly rely upon quasi-random and experimental research designs to identify peer effects empirically.

Research into these effects is especially prevalent in settings where randomization is common, such as classrooms and college dorms. This voluminous literature has demonstrated the impact of peer spillovers on a very wide array of outcomes ranging from education, income, physical fitness, risky behaviors, and attitudes towards minorities (Carrell et al., 2009, 2011, 2018, 2019; Chetty et al., 2011; Chung & Zou, 2023; Eisenberg et al., 2014; Feld & Zölitz, 2017; Hoxby, 2000; Kremer & Levy, 2008; Marmaros & Sacerdote, 2002; Sacerdote, 2001; Yakusheva et al., 2014). Another commonly studied setting is the workplace, where peer effects are often observed through learning, productivity, and behavioral influences through social pressure (Cohen-Zada et al., 2024; Cornelissen et al., 2017; Falk & Ichino, 2006; Guryan et al., 2009; Mas & Moretti, 2009; Oster & Thornton, 2012; Rosaz et al., 2016; Stevenson, 2017). A notable and related addition to this literature examined peer spillovers in effort and productivity in the Israeli Professional Football (soccer) Leagues (Cohen-Zada et al., 2024).

A large body of literature examines peer effects in financial matters, finding evidence of impacts on topics ranging from retirement savings to charitable giving (Duflo & Saez, 2003; Lieber & Skimmyhorn, 2018). Recent research also indicates that physical proximity to friends and colleagues is not the only channel through which peer effects occur. For instance, Bailey et al. (2018) and Hu (2022) show that geographically distant friends who experience either increases in the value of their home or exposure to major flooding events are more likely to purchase a home or flood insurance themselves. However, the most relevant strand of research to this paper focuses specifically on how peer health shocks influence behavior. For example, Innocenti et al. (2019) find that vicariously experiencing an acquaintance's negative health shock increases the intention to purchase health insurance more than the impact of one's illness. Similarly, Robertson et al. (1972) demonstrate that having a friend experience (but not die from) a car crash significantly increases seat belt usage through fear of injury.[1]

These findings from related literature lead me to investigate the hypothesis that individuals

[1]Interestingly, individuals seemingly perceive the probability of their death from a car accident as too improbable to cause changes in behavior whereas seeing a friend injured from a wreck was more salient.

who experience the CTE death of a former teammate will have differentially large labor supply responses relative to characteristically similar non-treated athletes. The first dimension of labor supply studied is at the extensive margin – whether to exit or remain in the profession. Research into the determinants of "early" retirement decisions has highlighted the complex nature of these decisions. Models that incorporate wide arrays of information on workers, such as their wage rate, the state of their finances, their eligibility for pensions, and their current health status, are vastly more predictive of labor supply decisions than models that analyze the impact of these factors separately (K. H. Anderson & Burkhauser, 1985; Giustinelli & Shapiro, 2024; Quinn, 1977). Relatedly, empirical work has demonstrated that extensive margin labor supply elasticities are very heterogeneous. some of the highlighted mechanisms driving differences in retirement are threats to finances and health (Brown, 2001; Brown et al., 2010; Coile & Levine, 2007) and differences in age and education (Coile & Levine, 2011).

The second dimension of labor supply examined in this paper is earnings. The theory of compensating wage differentials suggests that observed wage differences reflect monetary and non-monetary variations in the desirability of different types of work across time and space (Rosen, 1986). This theory assumes that labor supply can adjust to changes in job desirability, leading to extensive research on the intersection of workplace safety issues and worker bargaining power. When workers have limited understanding of workplace accident risks and face local labor markets with few alternative options, the risk of injury may be only partially reflected in wages (Bender & Mridha, 2011; Fishback & Kantor, 1992; Lavetti, 2020; Mridha & Khan, 2013). A related consequence of non-competitive labor markets is that they often lead to higher quit rates, with significant implications for both worker financial stability and aggregate productivity (Böckerman & Ilmakunnas, 2009, 2020; Cottini et al., 2011).

If market-level labor supply is inelastic due to the limited mobility of workers of non-transferable skills across industries, then increased unionization of workers should ameliorate these market inefficiencies via increased worker bargaining power. Research into the effect of unions on compensating differentials has shown that policies that relaxed the right for workers to unionize helped to increase wages and minimize the pass-through of costs of worker's compensation policies but were less successful in reducing the risks of work-related injuries (Fishback, 1986; Fishback & Kantor, 1995; Kim & Fishback, 1999). Importantly, the presence of compensating wage differentials is not dependent upon a high prevalence of workplace injuries. Hersch (1998) shows that there are large wage premia for women in white-color jobs who are exposed to small differences in unlikely work-related injury and illness.

Researchers have also long examined the consequences of these wage differences on inequality (Leeth & Ruser, 2003). For instance, Lavetti & Schmutte (2023) demonstrates that women and men sort within the labor market differently based on physical risk but similarly on financial risk, which contributes to establishment segregation and can explain a significant portion of the gender wage gap. Hersch (2011) provides evidence of another mechanism driving these differences in occupational segregation by demonstrating that female workers employed in settings with greater risks of sexual harassment earn more in wages, all else equal. While significant stressors such as performance pressure within the workplace have been shown to contribute meaningfully to inequality (Nagler et al., 2023), even small differences in preferences for seemingly innocuous factors such as commuting time and driving speed have been shown to have large impacts on aggregate gender wage gaps (Cook et al., 2021; Le Barbanchon et al., 2021).

Research on compensating wage differentials has also been examined in across various sports settings to provide unique insights into the topic. For example, Michaelides (2010) shows that the wages of professional players are highly elastic with respect to location amenities and non-pecuniary characteristics of the team. P. Anderson (2022) examines whether wages vary significantly to compensate workers for the risk of re-injury, finding large wage premia between professional boxers who have lost a fight via knockout. Additionally, Dole & Kassis (2010) finds that players with larger bargaining power (as measured by those in the top quartile of the income distribution) receive wage premiums for playing more games on surfaces commonly deemed to increase the probability of injury.

## 3 Setting

### 3.a. CTE & Collision Sports

Collision sports have long been recognized for their extreme health risks, including rare but severe incidents like paralysis and death, often resulting from blunt-force trauma to the spine, surgical complications, or serious lacerations. Although such worst-case outcomes were historically rare, the rising popularity of sports like American football in the United States, along with the advent of 24-hour sports news networks, brought increased attention to the health risks associated with these sports. A series of high-profile injuries in the 1990s, in which star NFL players were left unconscious during nationally televised games, propelled brain injuries into the national spotlight. The growing media attention on these well-known public figures experiencing often tragic cognitive declines made it increasingly difficult for the public to reconcile these declines with the athletes' once-impervious, superhuman personas.

Concussions, the temporary loss of brain function caused by a violent blow to the head, were a largely misunderstood injury, often deemed "invisible" due to their difficulty to detect through standard imaging techniques like MRI or CT scans, leaving diagnosis reliant on subjective symptoms and behavioral cues. Similar to how early misnomers related to HIV unintentionally downplayed the risks associated with contracting the virus (Black, 1986; Cardazzi et al., 2023; Shilts, 1987), dangerously misleading terminology was often employed to describe symptoms of head trauma in sports.[2] Similarly, athletes experiencing head trauma were commonly described as having their "bell rung" or being "punch drunk" (Martland, 1928), framing these injuries as rare, temporary, and typically affecting only reckless individuals. The metaphor of a "bell ringing" suggests that the effects would eventually stop, just as a drinker returns to sobriety after ceasing to drink, thereby obscuring the seriousness and permanence of these injuries. Thus, concussions were treated with the same level of seriousness as a sprained ankle – an injury that players could play through and recover from quickly if they had enough "toughness", and one thought unlikely to have significant long-term consequences (Fainaru-Wada & Fainaru, 2013).

This lack of clarity around what constituted a concussion, combined with the pressures of a sports environment where players, coaches, and team owners were all heavily incentivized towards risk, exacerbated and prolonged athletes' exposure to head trauma.[3] Scattered calls for reform of these sports were hampered by the lack of data on the prevalence or long-term consequences of these injuries.[4]

The tipping-point in the academic links between concussions and collision sports occurred in the mid-2000s following two landmark studies that identified chronic traumatic encephalopathy (CTE) in former NFL players (Omalu et al., 2005; Omalu et al., 2006). This discovery challenged prevailing assumptions that health risks in professional sports were limited to rare, catastrophic incidents. Instead, it revealed that repeated head impacts – common and often unavoidable in these sports – lead to severe and progressive neurological damage.

CTE is an irreversible neurodegenerative condition caused by repeated head impacts sustained over long periods (Ling et al., 2015). This disease is highly salient to athletes due to the severity

[2] In the 1980s, HIV was referred to as "GRID" (gay-related immune deficiency) and as a "4H disease" by medical professionals and the CDC, referencing those perceived to be most at risk: homosexuals, hemophiliacs, Haitians, and heroin users. This terminology led many to mistakenly believe that individuals outside these groups were not at risk.

[3] Players and coaches face a prisoner's dilemma regarding health and playing time: while all players would benefit from resting to recover, those who take the necessary time off risk of being replaced. Painkillers like Toradol further reduced players ability to "defect" from this cooperative equilibrium. Coaches encounter similar pressures to keep key players on the field, regardless of health. Meanwhile, team owners, reaping the rewards of soaring revenues and franchise valuations in sports with inherent dangers, faced strong incentives to downplay the public's growing safety concerns. These forces combined to create a strong culture of denial around concussions.

[4] See Fainaru-Wada & Fainaru (2013) for an overview of the NFL's pseudo-scientific, public relations entity which co-opted an academic journal to downplay the severity of concussions, the FTC's and academics' rebuking of claims published in these journals, and the class action lawsuits filed against the league for fraud and negligence for concealing the health risks of football.

of the symptoms and the relatively young age at which they can manifest. Individuals with CTE experience a wide range of impairments, including motor dysfunction, difficulties with emotional regulation, nerve pain, and, in many cases, dementia – symptoms that closely resemble advanced stages of Parkinson's and Alzheimer's disease. With no curative treatments available, care is limited to managing symptoms as the condition progressively worsens. Although these symptoms are more commonly associated with older populations, CTE is particularly concerning because it can begin affecting individuals as early as their mid-to-late 20s. Currently, there are no diagnostic tools to assess a person's risk of developing CTE as diagnosis can only be confirmed post-mortem through brain tissue analysis. While CTE can theoretically affect anyone, it has become synonymous with collision sports – particularly American football.

The discovery of CTE in former NFL players triggered a wave of diagnoses among other former athletes, all of whom shared similar traits: professional careers spanning a decade or more, disproportionate exposure to repeated head trauma, and aggressive playing styles. While the emergence of CTE among these athletes was undeniably disturbing, it initially implicated a relatively small group of their peers as being "at risk" for developing the disease, given the limited number of individuals with similar characteristics. Perceptions of who was at-risk began to change in 2009 with the discovery of CTE following the accidental death of Chris Henry. Unlike the previous cases, Henry was collision-averse and had played competitive football for fewer than ten years before his death at the age of 26. Then, within a 14 month span from 2010 to 2011, four hockey players – three of whom were still active professionals – died and were diagnosed with CTE at the ages of 45, 35, 28, and 27.[5]

Figure 1 presents two figures comparing the discovery of CTE among young athletes across two sports.[6] Figure 1a displays the timeline of CTE diagnoses among former professional athletes. In American football, diagnoses have been relatively evenly spaced, with annual fluctuations between 0 and 2 diagnoses annually for the past two decades. This contrasts sharply with hockey, which experienced a spike of 4 CTE diagnoses in 2010 and 2011, followed by relatively few in subsequent years. Figure 1b shows that, due to the relatively young ages of this first cohort diagnosed in hockey, there are significant differences in the immediacy of these deaths compared to those in football. Among the share of professional athletes who have ever had a former teammate die of CTE, hockey players are twenty percentage points more likely to have experienced this health shock less than one year after they were last teammates.

[5]Notably, each player was recognized as an "enforcer" – a role designated to deter and retaliate against dirty or violent conduct from the opposing team, often through aggressive tactics which exposed them to a disproportionate number of blows to the head per game. However, the physical toll of enforcing and as these players' relative lack of skill means there are neither many enforcers nor do they have particularly long careers.

[6]By "young athletes", I specifically refer to those with at least five former teammates still playing professionally at least one year before their death.

**Figure 1.** Differences in the Discovery of CTE among Young Athletes across Sports

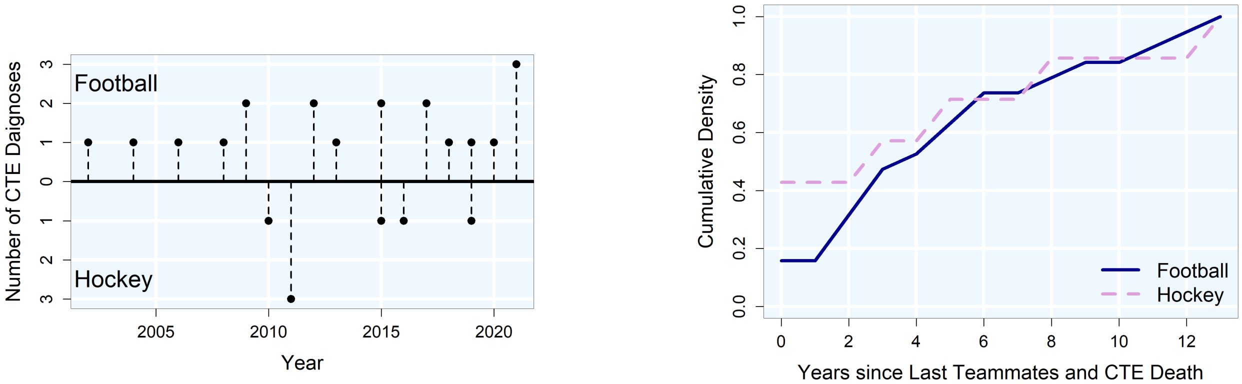


**(a)** Timeline of CTE Deaths

**(b)** Density of Recent CTE Deaths

The CTE-related deaths of hockey players came as a shock and brought the disease to the forefront of public awareness regarding the dangers of collision sports. Despite the inherently violent, collision-heavy nature of hockey – which involves players often weighing over 200 pounds colliding at speeds exceeding 20 miles per hour – hockey was thought to carry a significantly lower risk of CTE compared to American football. This belief likely stemmed from the lower incidence of high-profile concussions and instances of players being knocked unconscious during games. Additionally, the sport's culture, which emphasized aggression as a key strategy for success and privileged toughness over worker safety, further contributed to downplaying the risks. However, these deaths shifted the narrative from focusing solely on concussions to recognizing the dangers of repeated blows to the head and neck.

The inability to diagnose CTE, coupled with likely strong endogeneity of testing, means that there is no reliable estimate of the prevalence of the disease. This has lead to a tremendous variation in estimated risk. For example, researchers at Boston University's Chronic Traumatic Encephalopathy Center revealed that over 90% of nearly 400 brains of former athletes who were studied had signs of CTE (Chobanian & Avedisian School of Medicine, 2023). Relatedly, a recent survey of over 4000 former professional football players revealed than over 33% believe they are currently living with the disease (Grashow et al., 2024).

### 3.b. Labor Markets for Professional Athletes

North American sports leagues have implemented various mechanisms that shape the labor markets for professional athletes through collective bargaining agreements (CBAs) with the respective Players' Assocations. In both the NHL and NFL, CBAs enforce strict salary caps, which limit the total amount teams can spend on player salaries each year. They also establish restricted free agency, limiting players' ability to negotiate contract terms with other teams after their contracts expire, and impose numerous price controls, including maximum salaries, rookie salary caps, and veteran minimums. This system collectively limits players' bargaining power and ability to fully capitalize on their talents, often at the height of their careers.

While these constraints on earnings certainly limit workers' total earnings and variation within income bins, wages for athletes are still determined via a competitive bidding process. Figure 2 provides evidence of this hypothesis, plotting average player productivity and the average of their wages, revealing a strikingly positive linear trend for both sports. In addition to being paid according to one's ability to "produce" wins in sports, athletes are compensated according to the demand for the product they produce.[7] Figure 2 reveals large differences in the financial returns to productivity across the two sports with a one standard deviation increase in productivity increasing wages for hockey players by $1M (USD) and $2.5M (USD) for football players.

**Figure 2.** Productivity and Salary

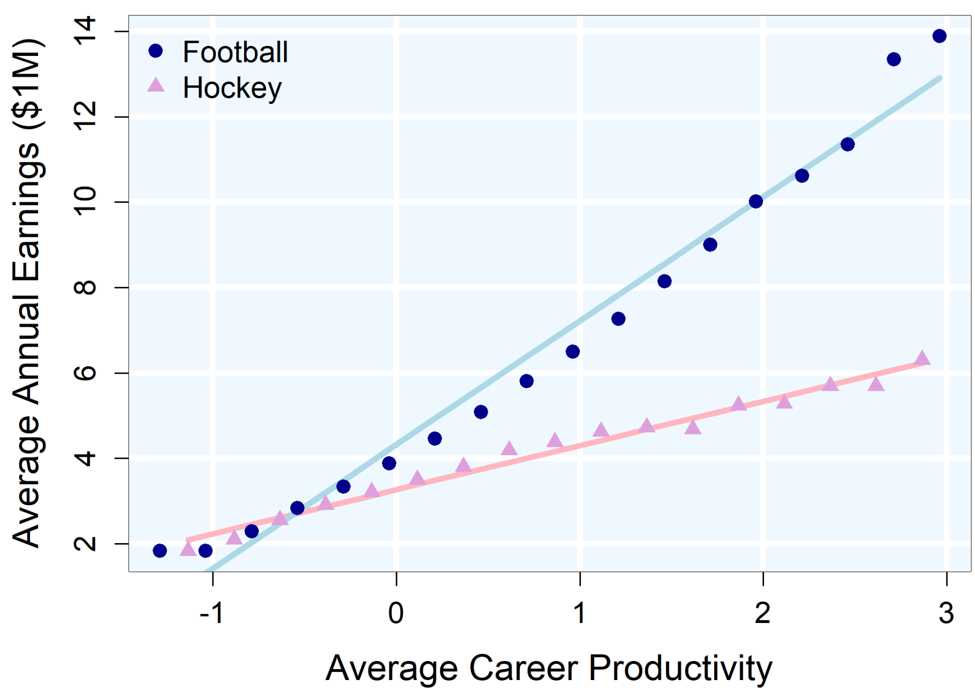


Figure 2 reveals a notable "superstar" effect in the labor market, with clear asymmetric returns at the top of the ability distribution.[8] Interestingly, this effect appears only among football play-

[7]In 2024, the salary cap per worker was $6.2M (USD) in American football and $3.8M (USD) in hockey. Further, figure A5 in the Appendix shows that the championship game for American football regularly draws 20 to 40 times more viewers than the championship series in hockey, reflecting the large disparities in revenue generated by each league.

[8]For a more in-depth theoretical exploration of the economics of superstars, see Rosen (1981) and Adler

ers. Two main factors likely contribute to this phenomenon. First, the rules of football promote superstar dynamics by offering greater space per player and favoring offensive playmaking, creating conditions conducive to standout performance.[9] Second, football players are able to specialize in position-specific skills more extensively than hockey players, largely due to the structure and demand of the game. Frequent pauses allow for a natural rotation of players, and larger roster sizes – driven by high demand for participation and spectator interest – enable players to focus on specific positions in offense or defense. In contrast, hockey skaters must be more adaptable, often switching between offensive and defensive roles as needed.

Furthermore, wage variation for hockey players is far more limited than in football, largely due to provisions in the NHL's CBA. Signed in 2005 during a period of inconsistent demand for the sport, the CBA secured "fully" guaranteed contracts for players in exchange for several concessions to owners.[10] Among these concessions, most notably, were stricter salary cap regulations and the possibility of escrow payments – meaning a portion of player salaries could be withheld if league revenues fall below an agreed – upon threshold. This arrangement contrasts sharply with American football, where a large pool of replacement talent diminishes union bargaining power, limiting athletes' ability to secure fully guaranteed contracts. Consequently, employment protections in football are minimal, with guaranteed compensation primarily serving as an incentive to attract highly productive free agents.

This creates a stark contrast in the financial security athletes experience upon exit across different sports. In hockey, if a team terminates a player's contract, the team must still pay the remaining guaranteed salary as stipulated in the contract, providing athletes with a significant level of financial protection. In football, however, players lack this security; if their contracts are terminated, they forfeit any non-guaranteed portions of their salary, often leaving them without the financial stability enjoyed by athletes in other sports.

## 4 Theoretical Framework

I formalize a mathematical framework of how workers' labor supply decisions can be influenced by the salience of peer health shocks in the context of violent, collision sports. Workers maximize lifetime utility by choosing whether to exit the profession each period, balancing the trade-offs between wage differentials across professions with greater risks to one's health and longevity.

(1985).

[9]See figure **??** in the Appendix for a greater discussion of the distribution of productivity across sports.

[10]In the NHL, guaranteed compensation applies primarily to players who have reached a certain age or have accrued a specific number of years of experience, typically after three years on entry-level contracts. After this point, players benefit from guaranteed contracts as long as they are not bought out, which involves compensation at a reduced rate.

This is expressed in Equation 1,

$$U_i = \max_{d_{ist}} \sum_{t=T_0}^{T} \delta^{t-T_0} \left( d_{ist} \cdot \omega_{ist} + (1 - d_{ist}) \cdot \eta_{ist} \right) \tag{1}$$

where workers ($i$) maximize utility ($U$) from the present period ($t$) until retirement ($T$) by deciding ($d \in \{0, 1\}$) whether to continue earning wages from professional sports ($\omega$) or exit to earn wages in their next best alternative profession ($\eta$). The discount factor ($\delta$), reflects time preferences, weighting future wages less than current wages.

The constraining trade-off in this utility maximization problem is shown in Equation 2, which illustrates that non-sport wages are determined by two factors: the worker's human capital in the labor market ($\mu$) and the number of years they can work.[11] Specifically, in this context, continued participation in professional sports is physically taxing, which reduces the number of years a worker can earn non-sport wages by some amount ($\theta$).[12] I will broadly refer to this parameter as the longevity penalty.

$$\eta_{ist} = \mu_{it} - \theta_{ist} \tag{2}$$

I allow the longevity penalty to vary across individuals, as shown in Equation 3. The penalty consists of two components: known workplace hazards ($\gamma$) and a residual term ($\epsilon$). The first component reflects the risks of workplace participation known to all athletes within each sport by year, while the residual captures all remaining individual-specific disamenities. Common theoretical assumptions – whether workers have complete information about workplace safety or that residual disamenities are randomly distributed – predict that the expected value of the residual term is zero.

$$\theta_{ist} = \gamma_{st} + \epsilon_{ist} \tag{3}$$

Equation 4 illustrates that the second major component of worker utility, sport wages, is determined by two primary factors. The first factor represents the value of the marginal product of labor through the product of market demand ($\rho$) and worker's productivity ($\upsilon$).[13] Second, the

[11] It is assumed that one cannot improve their human capital while employed as a professional athlete, so their human capital is fixed prior to period $t$.

[12] For simplicity, this parameter is modeled as a reduction in annual non-sport wages, rather than in the number of years one can work ($T$).

[13] An implicit assumption of this equation is that market demand for a sport is not causally related to the risks its workers face. While this assumption is clearly false in a literal sense, the highly inelastic demand for both

theory of compensating wage differentials predicts that wages are influenced by the disamenity associated with participating in the sport ($\theta$) which manifests in Equation 3 as long-term health issues. The degree to which these disamenities are reflected in wages depends on worker's bargaining power ($\beta$).

$$\omega_{ist} = \rho_{st} \cdot \upsilon_{ist} + \beta_{st} \cdot \theta_{ist} \tag{4}$$

The factors influencing worker bargaining power are shown in equation 5. Game theoretically, the credibility of player's threats improve bargaining outcomes. The credibility of the threat to retire from professional sports is captured by the first parameter which measures the value of the worker's human capital in the labor market ($\mu$). The second component in this equation measures superstar effects. There are imperfect substitutes for highly skilled labor. The number of suitable replacement workers shrinks as worker productivity ($\upsilon$) increases, which increases worker bargaining power. This measure of productivity is augmented with a measure of the strength of the collective bargaining agreement ($\pi$) binding the worker by sport and year.

$$\beta_{ist} = \mu_i + \pi_{st} \cdot \upsilon_{ist} \tag{5}$$

Very simply, each period workers will remain working in sports if their sport wage is higher than their non-sport wage ($\omega_{ist} \geq \eta_{ist}$). Representing this decision by substituting equations 2 and 2 can be seen in 6.

$$\rho_{st} \cdot \upsilon_{ist} + (\mu_i + \pi_{st} \cdot \upsilon_{ist}) \cdot (\gamma_{st} + \epsilon_{ist}) \geq \mu_i - (\gamma_{st} + \epsilon_{ist}) \tag{6}$$

This model can be simplified by incorporating unit and time fixed effects, helping to isolate relevant factors related to what cause *changes* in individual labor supply decisions. These fixed effects differences out factors that are constant over time for each individual and for factors that are constant across individuals within the same time period.[14] The results of this process are displayed in Equation 7.

hockey and football displayed in Figure A5 with respect to the wide-spread health concerns for the athletes at the time suggests that any relationship between the two variables is likely to provide both little explanatory benefit and possibly even obfuscate the model through additional complexity.

[14]Unit fixed-effects can be mathematically expressed as $\Delta X_i = X_{it} - \frac{1}{T}\sum_{\tau=1}^{T} x_{i\tau}$ while time fixed effects can be written as $\Delta X_i = X_{it} - \frac{1}{I}\sum_{\iota=1}^{I} x_{\iota t}$. For notational simplicity, I denote unit and time fixed-effects as $\Delta_i$ and $\Delta_t$ respectively.

$$\Delta_{it}(\upsilon_{ist} + \epsilon_{ist} \cdot (\upsilon_{ist} + 1)) \geq 0 \tag{7}$$

The results from Equation 7 yield two theoretical predictions. First, if workers have full information about workplace safety, or if residual risk perceptions do not vary systematically across workers, changes in productivity are the only relevant factor explaining voluntary exits from risky workplaces. Second, if increased risk perceptions affect different groups of workers unevenly, the impact of such shocks on exit decisions will depend on whether a worker's productivity is rising or falling. Though this outcome may seem counterintuitive, it is theoretically consistent. For example, more productive workers can expect greater gains from remaining in the sport after a negative health shock, as their higher productivity increases future wages both directly, through the value of their marginal product and bargaining power, and indirectly, through the exit of less productive workers. Conversely, workers with declining productivity will be more likely to exit following a work-related health shock for the same reasons.

## 5 Data

I gather data from two sources to empirically test the impact of peer health shocks on labor supply. The primary dataset is sourced from Sports Reference, a collection of websites offering comprehensive statistics and historical data across various sports. A key advantage of this dataset is its wealth of athlete-specific information, which allows me to track the movements of these workers both leading up to and throughout their professional careers. Importantly, this enables me to observe not only their former teammates during their professional careers but also occasionally during college, semi-professional teams, and even high school.

For analysis, I construct player-by-season records for all professional American football and hockey players who made the final roster for any team in the premier divisions of these sports from 1970 to 2023. The profiles of these athletes include a variety of details, such as years of participation in university or amateur teams, the name of their high school, city of birth, athlete relatives, recruiting rank, draft order, date of birth, and, when applicable, date of death. Additionally, I gather data on positional information, accolades, preseason championship expectations (as measured by implied probabilities from pre-season betting markets), and game participation.

Most importantly, for identification purposes, this dataset contains continuous measures of worker productivity. For hockey, I use player-by-season estimates of "Point Shares," which attribute a player's contributions to their team's total points in a season. In football, I employ the "Approximate Value" metric, which similarly measures a player's contribution to their team's chances of winning.[15] These measures are notable for their ability to harmonize statistics across players in vastly different positions within their sport, while also making productivity comparable across time, even as the rules and strategies of the sports evolve.[16]

To account for differences in the number of games played by teams over time, I standardize these productivity measures within each sport and year, such that a one-unit change reflects a one standard deviation difference in productivity relative to other athletes in the same sport and year. As a result, the mean value of productivity in this analysis is zero, which reflects the threshold below which athletes are likely to be fired and replaced.[17]

I then match this with a secondary dataset sourced from Spotrac, a website that provides financial information about player contracts in various professional sports, including American football and hockey. Though Spotrac is primarily focused on tracking salary caps – that is, the extent to which teams are bound by league-imposed spending limits on player salaries, this platform offers the most detailed and comprehensive data on contracts signed by professional athletes available at the time of this writing. The dataset includes information on the total value of contracts, the duration of employment, and detailed breakdowns of compensation into guaranteed and variable components over time. Additionally, contracts are dynamically updated to reflect extensions and terminations. This rich dataset allows me to calculate the exact dollar amount workers would forgo upon retirement in each year.

After merging these datasets, the final analysis sample consists of 2,312 unique hockey players, of whom 977 (42%) have complete salary information, and 538 (23%) are identified as having had a former teammate die from CTE. The sample also includes 8,817 unique football players, with 4,916 (56%) having complete salary data and 1,244 (14%) are identified as having had a former teammate die from CTE.[18]

## 6 Identification

To examine the differential impact of health shocks on peer labor supply decisions in professional sports, I compile a list of all former athletes reported to have been diagnosed with CTE. Tracking these athletes across their professional, semi-professional, and occasionally high school careers

[15]These measures are slight derivations of the concept of "Win Shares," introduced in the late 1970s by writer and statistician Bill James, who, perhaps unsurprisingly, holds a degree in Economics as these measures often closely resemble concepts from the discipline such as the marginal product of labor.

[16]For a detailed analysis of the correlates of worker productivity, see Figure A1 in the Appendix.

[17]For an in-depth explanation of how Approximate Value and Point Shares are calculated, see pro-football-reference.com/about/approximate_value.htm and hockey-reference.com/about/point_shares.html.

[18]For more information regarding missingness of contracts over time and across players, see Figure A10 in the Appendix.

enables me to construct a social network matrix of athletes who have played professional hockey and football.[19] This social network matrix allows me to identify a near-complete list of all individuals who were ever teammates with someone who died of CTE, revealing how long they were teammates and how much time has passed since they last shared a team.

A list of former athletes who have been diagnosed with CTE is presented in Table 1, including only those with at least 10 former teammates still active professionals four years before their death for visual simplicity. This table presents the primary source of identifying variation used in this paper—the date of death of these athletes. The selection criteria require that each athlete had at least 30 former teammates still active in professional sports four years before their death, with at least five still active one year before their death.[20] A sufficiently large sample of treated workers is required for each treatment event to ensure that the results are robust and not overly sensitive to random fluctuations in smaller sample sizes.

I conduct a difference-in-differences analysis, comparing the labor supply choices of characteristically similar athletes who differ only in their exposure to a colleague who died of CTE, before and after the time of the death of their former teammate. This setting provides unique advantages for causal inference, such as limited opportunities for workers to self-select into teams and the quasi-random timing of CTE-related deaths. However, I cannot estimate causal effects directly using this natural experiment, as news of CTE diagnoses for former players reaches everyone—both peers and non-peers.[21] As a result, this empirical approach measures the *differential* impact of these deaths on former teammates. Consequently, if these deaths influence labor supply equally across all athletes, the estimates from these comparisons should, in expectation, be equal to zero.

A key assumption in difference-in-differences models is that, in the absence of treatment, trends in the outcome would evolve in parallel between comparison groups over time. However, this assumption is unlikely to hold in this setting due to high turnover rates in professional sports. For instance, annual turnover is roughly 20% in football and 14% in hockey, as shown in Figure A3a, resulting in notable differences in career length – the median career spans four years in hockey compared to three years in football, as shown in Figure A3b. If being "treated" by a former colleague's death requires that a player was once teammates with an athlete who has since passed away, then treated workers are likely to be significantly older and more productive than typical non- or not-yet-treated workers. The greater the time lag between an individual's re-

[19] In this context, "semi-professional" often refers to collegiate careers. For football players, this is almost always the case, while for hockey players, it includes collegiate as well as amateur, international, and lower division teams.

[20] All other athletes treated by the death of a former colleague are omitted from the analysis.

[21] Figure A6 in the Appendix displays time-series variation in search intensity for "CTE", frequently peaking at the timing of CTE deaths and diagnoses.

**Table 1.** CTE-Diagnosed Athletes with Remaining Teammates at Time of Death

| Player | Sport | Career | Date of Death | Teammates at $\tau = -4$ |
|---|---|---|---|---|
| Mike Webster* | Football | 1974-1990 | 9/24/2002 | 22 |
| Justin Strzelczyk | Football | 1990-1998 | 9/30/2004 | 61 |
| Andre Waters* | Football | 1984-1995 | 11/20/2006 | 13 |
| Tom McHale* | Football | 1987-1995 | 5/25/2008 | 16 |
| Shane Dronett | Football | 1992-2001 | 1/21/2009 | 40 |
| Chris Henry | Football | 2005-2009 | 12/17/2009 | 137 |
| Bob Probert | Hockey | 1986-2002 | 7/5/2010 | 53 |
| Derek Boogaard | Hockey | 2006-2011 | 5/13/2011 | 102 |
| Rick Rypien | Hockey | 2006-2011 | 8/15/2011 | 92 |
| Wade Belak | Hockey | 1997-2011 | 8/31/2011 | 151 |
| Junior Seau | Football | 1990-2012 | 5/2/2012 | 157 |
| Jovan Belcher | Football | 2009-2012 | 12/1/2012 | 116 |
| Paul Oliver | Football | 2008-2011 | 9/24/2013 | 115 |
| Steve Montador | Hockey | 2002-2012 | 2/15/2015 | 141 |
| Adrian Robinson | Football | 2012-2013 | 5/16/2015 | 115 |
| Tyler Sash | Football | 2011-2012 | 9/8/2015 | 80 |
| Marek Svatoš | Hockey | 2004-2011 | 11/4/2016 | 59 |
| Aaron Hernandez | Football | 2010-2012 | 4/19/2017 | 71 |
| Daniel Te'o-Nesheim | Football | 2010-2013 | 10/29/2017 | 117 |
| Kevin Ellison* | Football | 2009-2009 | 10/4/2018 | 16 |
| Greg Johnson* | Hockey | 1994-2006 | 7/7/2019 | 14 |
| George Atkinson | Football | 2014-2016 | 12/2/2019 | 117 |
| Max Tuerk | Football | 2017-2017 | 6/20/2020 | 68 |
| Vincent Jackson | Football | 2005-2016 | 2/15/2021 | 92 |
| Phillip Adams | Football | 2010-2015 | 4/8/2021 | 114 |
| Demaryius Thomas | Football | 2010-2019 | 12/9/2021 | 199 |

This table represents all individuals who were diagnosed with CTE and had at least ten former teammates playing professional sports at least four years prior to their death. Star symbols (*) are used to indicate those whose teammates are omitted from the analysis.

tirement and their eventual CTE-related death, the larger the disparity in characteristics between their former teammates and the remaining workforce are likely to be. As these differences grow, the ability to attribute observed differences in outcome variables to the salience of work-related disamenities weakens. Visual confirmation of concerns that using the full-sample of non- and not-yet-treated workers as a control group is provided in Figure A4 which shows substantial differences in exit rates between treated and untreated workers leading up to the CTE death of a former teammate.

To address concerns over differences between comparison groups, I employ a Coarsened Exact Matching (CEM) method, which adjusts for pre-treatment covariate differences between treated and control groups. The primary advantage of CEM over other matching methods is that it ensures both the mean and distribution of each covariate are similar across groups by "coarsening" their values into small bins, or "strata." This method generates two key outputs: indicators for poorly matched control units and a numerical weight that accounts for any remaining differences between the groups. I primarily rely on the former output, as the large number of workers in both the treated and control groups often makes weighting unnecessary.

The results from this matching process, which relies solely on two covariates, worker age and productivity, are displayed in Figure 3. As expected, treated workers differ significantly from the full sample control group, earning approximately 13% higher salaries, being 5% older, and demonstrating 10% greater productivity. After matching, however, no statistically significant differences remain, as the process excludes control group athletes who are the most dissimilar from the treated group in these characteristics.

Figure 4 demonstrates that, in the matched sample, the pre-treatment trends in the outcome variable are both parallel and approximately equivalent across the two groups.[22] The figure also shows a significant increase in the rate of exit for treated workers, which occurs only after the timing of the treatment, suggesting a significant treatment effect.

### 6.a. Endogenous Bahavioral Responses

While the models estimated with individual-level data help mitigate concerns about omitted variable bias in treatment effects, identification may still be compromised by endogenous behavioral responses to CTE deaths, which disproportionately impact treated players over time. For instance, unobserved individual or team-level responses—such as increased shirking or greater use of per-

[22]This is particularly important given the bounded nature of the outcome variable. For example, the rate of cohort exits is likely to asymptotically approach an upper bound of one over time. This suggests that groups with lower exit rates prior to treatment may close gaps in the outcome due to inherent properties of the variable, rather than treatment effects.

**Figure 3.** Covariate Balance

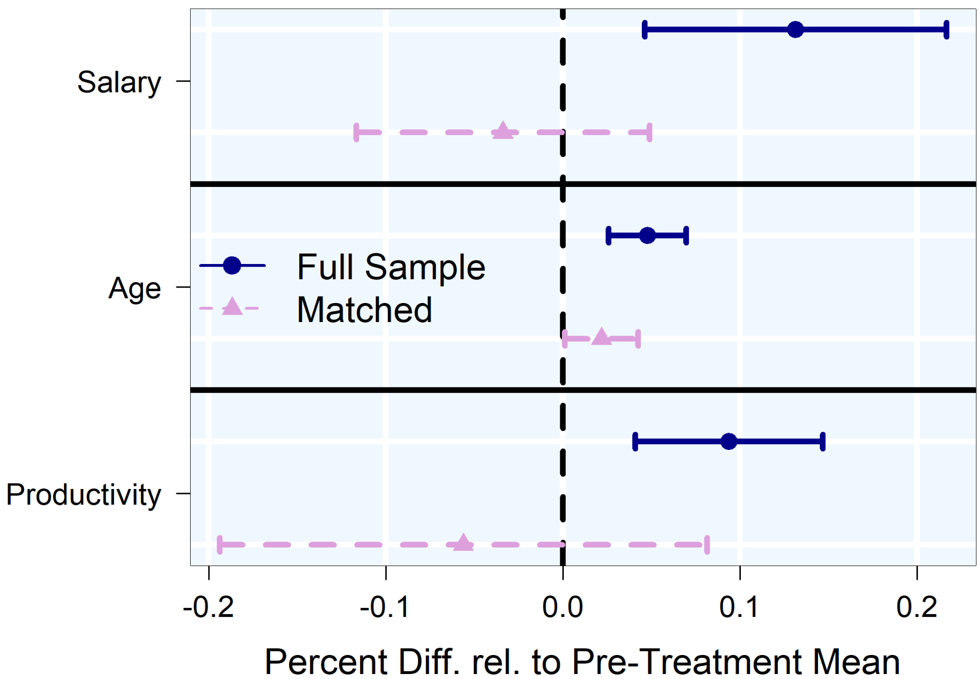


**Figure 4.** Matched Trends in Cumulative Share of Cohort Exits

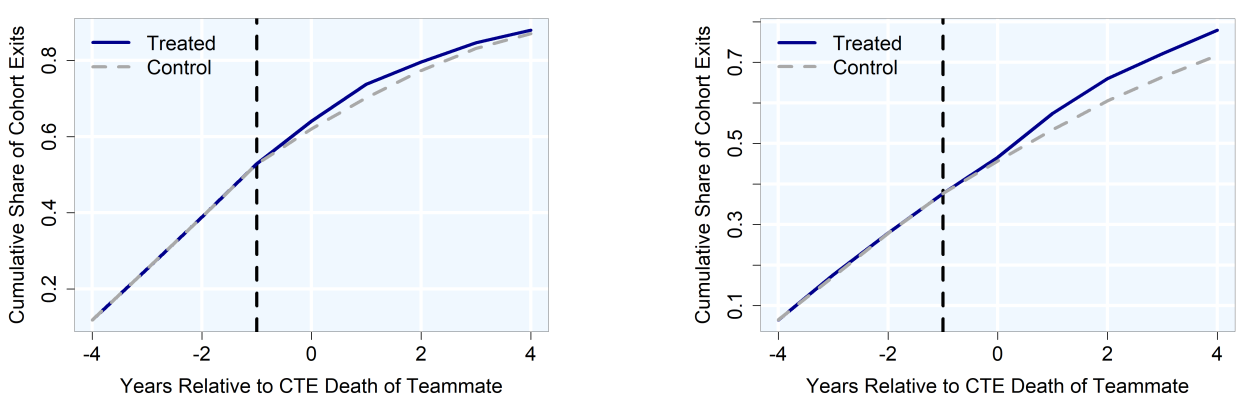


**(a)** Football: Matching

**(b)** Hockey: Matching

sonal protective equipment—could bias estimates of the 'true' treatment effect, as players would have less motivation to alter their behavior in the absence of heightened injury risk.

There are several key points in this sports context that help alleviate concerns about threats to identification. First, the wages earned by these athletes are significantly higher than what they would likely earn outside of professional sports, creating a large surplus of workers ready to replace any athlete who shirks their duties. Additionally, coaches exogenously determine players' labor supply decisions at the intensive margin. Failure to follow orders from team management provides legal grounds for teams to terminate players' contracts.[23] Furthermore, professional sports are highly competitive, with player productivity as the central focus. Teams are quick to release players whose marginal value falls below the marginal cost of employment.

Second, most safety policies and equipment changes do little to mitigate the risk of developing CTE. As team sports have become more professionalized with greater emphasis on health and recovery, athletes (relatively) rarely sustain concussions during practice. The highest-risk contact occurs during games, where opponents are incentivized to use forceful, high-impact maneuvers. Attempts to limit concussions in collision sports are likely ineffective, as helmets and other protective equipment primarily aid in impact absorption rather than preventing brain movement within the skull. Meaningful reductions in concussions, similar to interventions in auto racing, would require a firm neck restraint to prevent whiplash on contact, a solution that is not feasible for sports like football (Kaul et al., 2016). Nevertheless, players' beliefs in the efficacy of these policies ultimately influence their behavior.[24] Even if all treated players fully believed in the effectiveness of these interventions, any resulting bias in treatment effect estimates would likely bias treatment effect estimates towards zero rather than away from it.

## 7 Estimation

I next employ difference-in-differences models, which more convincingly isolate the effects of treatment through the use of worker and year fixed effects. Time-invariant individual characteristics, such as players' height, race, education, position, draft order, and baseline talent, are accounted for with player fixed effects. Year fixed effects, in turn, capture period-specific changes in the dependent variable, such as labor strikes, adjustments in player safety rules or their enforcement, new collective bargaining agreements, fluctuations in league revenue, and the establishment of commissions to provide payouts for former athletes affected by work-related neurotrauma.

[23]Insubordinate athletes would be correctly classified in the model as having exited.

[24]However, this belief is unlikely, given the well-documented unintended effect that safety measures for head injuries often lead to increases in lower-body injuries (Hanson et al., 2017).

To address the well-known biases associated with using two-way fixed effects (TWFE) models with staggered treatment timing, as identified in the modern econometrics literature, I use a stacked event-time approach by treatment cohort. Commonly referred to as a "stacked" difference-in-difference estimator, this method follows the familiar TWFE estimation process but requires restructuring the data to center it in event-time. Specifically, stacking involves selecting an event-time window of nine years (four years pre-treatment and five years post-treatment for this analysis). After keeping only the information within this period, I store the IDs of all treated workers to ensure they are not later used as control units. Finally, I create a unique "stack" identifier for each cohort. This process is repeated for all unique treatment timings, and the stacks are ultimately combined into one dataframe for analysis.

The process of restructuring the data this way can be seen in the estimating equations below. Equation 8 tests worker $i$ in stack $s$ in year ($t$) exits ($E$) from professional sports. This model relies upon worker fixed effects ($\alpha$) and year fixed effects ($\delta$). A treatment indicator $T$ is equal to zero for all units prior to treatment, and then is set equal to one for all former teammates of a worker who died of CTE after their death. The coefficient $\beta$ represents the differential impact of the likelihood of exiting the profession after treatment between treated and control groups.

$$E_{sit} = \alpha_{si} + \delta_{st} + \beta T_{sit} + \epsilon_{sit} \tag{8}$$

Equation 9 estimates the dynamics of the effect of a peer health shock on the probability of exit. The period immediately prior to treatment ($\tau = -1$) serves as the reference period. There are two major advantages of this model. First, it allows one to test whether the assumption that the outcomes of the two groups would continue trending parallel to one another in the absence of treatment. Second, it allows one to visualize how the effects evolve over time.

$$E_{sit} = \alpha_{si} + \delta_{st} + \sum_{\tau=-5}^{-2} \beta_\tau T_{i\tau} + \sum_{\tau=0}^{-5} \gamma_\tau T_{i\tau} + \epsilon_{sit} \tag{9}$$

Equation 10 displays a Poisson Regression model which mirrors equation 8 but for its log transformation of the expected value of the dependent variable. This is done primarily for the salary outcomes which are highly skewed and in order to avoid concerns about the sensitivity of results which analyze log-transformed outcome variables (Chen & Roth, 2024). Thus,the coefficient of interest in this regression ($\beta$) measures the percentage (rather than percentage point) change in the dependent variable caused by treatment.

$$log(E[Y_{sit}]) = \alpha_{si} + \delta_{st} + \beta T_{sit} + \epsilon_{sit} \tag{10}$$

Last, this study employs both balanced panel and unbalanced data. The latter dataset can be conceptualized as a panel of team rosters, which contain detailed information about each worker. Crucially, no new workers are allowed to enter the sample after the treatment timing within each stack. The dependent variable of interest, "exit," is equal to zero until it takes a value of one the last time a worker is observed in the dataset. Crucially, this setup captures changes in the composition of the workforce over time. As such, only the remaining workers in each period serve as comparison groups for one another. This contrasts with the balanced panel dataset, which requires that individuals remain in the sample for every period. In this case, the dependent variable "exit" is set to zero until it takes a value of one for every period after the worker is last observed in the unbalanced panel.

Therefore, differences in estimated treatment effects between these groups can be interpreted as differences in comparison groups: either relative to all workers in the balanced panel or relative to only the remaining workers in each period within the unbalanced panel. This suggests that treatment effects may be more sensitive when estimated using the unbalanced data, where the gradual attrition of workers creates a more dynamic workforce composition, compared to the balanced panel, where comparison groups remain static over time. However, this dynamically changing workforce composition may bring the benefit of potential for more appropriate comparisons if there is endogenous in pre-treatment retirement decisions.

# 8 Results

## 8.a. Professional Exit (Retirement)

Table 2 presents empirical estimates of the differential impact of well-publicized deaths related to workplace safety on the probability that their former coworkers exit the profession. Estimating models using data from the entire universe of athletes on NFL and NHL rosters in the years surrounding these deaths, each regression model in Table 2 reveals that CTE deaths significantly increase the probability that former teammates retire.

The binary dependent variable used in Table 2, *Exit*, is multiplied by one hundred to express the coefficients in percentage terms, making the result easier to interpret and preventing the effect sizes from being obscured by small decimal values. Measures of worker productivity are standardized in order to provide a more easily comparable method for comparing the magnitude of the treatment effect on retirement with that of a one-standard-deviation change in worker

**Table 2.** The Effect of the CTE Death of a Former Teammate on Retirement

| Dependent Variable: | $\mathbb{1}$(Exit)*100 | | | |
|---|---|---|---|---|
| Sample: | Full | | Matched | |
| Model: | (1) | (2) | (3) | (4) |
| **Panel A (Balanced)** | | | | |
| CTE (Post) ×Teammate | 4.58*** | 1.79** | 2.53*** | 2.24*** |
| | (0.78) | (0.70) | (0.80) | (0.72) |
| Productivity | | -13.21*** | | -13.30*** |
| | | (0.08) | | (0.12) |
| Pre-Treatment Mean Retire | 29.03 | 29.03 | 29.03 | |
| % Change | 0.16 | 0.06 | 0.09 | 0.08 |
| Observations | 409,383 | 409,383 | 229,833 | 229,833 |
| **Panel B (Unbalanced)** | | | | |
| CTE (Post) ×Teammate | 3.03*** | 1.89* | 4.00*** | 3.97*** |
| | (1.05) | (1.00) | (1.09) | (1.04) |
| Productivity | | -10.03*** | | -10.12*** |
| | | (0.12) | | (0.16) |
| Pre-Treatment Mean | 17.28 | 17.28 | 17.28 | 17.28 |
| % Change | 0.18 | 0.11 | 0.23 | 0.23 |
| Observations | 190,886 | 190,797 | 110,561 | 110,516 |

Ordinary Least Squares (OLS) is used to estimate each model which all include stack-by-player and stack-by-year fixed-effects. The depednent variable *Exit* is a binary estimate that is multiplied by one hundred for presentional simplicity. The mean of the dependent variable is calculated from pre-treatment observations of players in the treated group. *Clustered (Stack-by-Player) standard-errors in parentheses. Signif. Codes: ***: 0.01, **: 0.05, *: 0.1.*

productivity.

Panel A of this table analyzes a balanced sample, meaning individuals do not drop out of the dataset after retirement. Thus, the coefficients in this panel represent changes and persistence in the probability of exiting due to treatment compared to all other workers on rosters prior to treatment. Panel B is unbalanced, meaning that coefficients on treatment indicators from this panel reflect changes in the probability of exiting relative to *remaining* workers. Estimates from the later model are inherently more sensitive due to the smaller sample sizes. However, they are included in the table because groups of remaining workers could reflect more suitable counterfactual groups in the periods surrounding treatment.

Models one and two of Table 2 are estimated using data on all workers. Treatment effect estimates across model one in each panel reflect significant treatment effects, which diminish significantly when conditioning on worker productivity. As discussed in the Identification section of this paper, the large change in these estimates from models one to two suggests significant differences in productivity across workers in the treatment and control groups. Thus, my preferred estimates come from models three and four of Table 2 which are estimated using data from a sample of workers who are characteristically similar prior to treatment. Adding a control for worker productivity in model four does not significantly alter the coefficients on the treatment indicators, suggesting the usage of the more appropriate control group, and to stress that the results are not sensitive to the choice of sample that is analyzed.

The results from model four of panels A and B suggest that treatment increases the probability of exit from the profession by 2.24-3.97 percentage points. Along with being statistically significant, these effect sizes are meaningfully large. The magnitude of the treatment indicators suggests that the CTE death of a former teammate increases the probability of retirement by 8-23% relative to the pre-treatment mean. This effect size is approximately equivalent to a 0.16-0.39 standard deviation reduction in worker productivity.

To explore the mechanisms driving this increase in the probability of retiring, I stratify these results by sport in Table 3. Models one and two (three and four) examine the treatment effects for hockey (football), while odd-numbered models use the full control group and even-numbered models use the matched samples. The key finding of this table is that the treatment effect is unambiguously larger in hockey than in football.

The results of Table 3 align with expectations, reflecting important contextual differences across sports, especially if familiarity and proximity are plausible mechanisms driving this outcome. For instance, the smaller team sizes, longer careers, extended periods as teammates, more abrupt CTE-related deaths, and the greater relative 'surprise' surrounding the discovery of CTE

**Table 3.** Peer Work-Related Deaths & Retirement: Heterogeneity by Sport

| Dependent Variable: | $\mathbb{1}$(Exit)*100 | | | |
|---|---|---|---|---|
| Sport: | Hockey | | Football | |
| Sample: | Full | Matched | Full | Matched |
| Model: | (1) | (2) | (3) | (4) |
| **Panel A (Balanced)** | | | | |
| CTE (Post) ×Teammate | 7.08*** | 4.42*** | 3.44*** | 1.61* |
| | (1.39) | (1.41) | (0.94) | (0.98) |
| Pre-Treatment Mean | 22.4 | 22.4 | 32.21 | 32.21 |
| % Change | 0.32 | 0.2 | 0.11 | 0.05 |
| Observations | 74,619 | 52,893 | 334,764 | 176,940 |
| **Panel B (Unbalanced)** | | | | |
| CTE (Post) ×Teammate | 5.08*** | 5.15*** | 1.76 | 3.21** |
| | (1.51) | (1.53) | (1.42) | (1.51) |
| Pre-Treatment Mean | 13.46 | 13.46 | 19.13 | 19.13 |
| % Change | 0.38 | 0.38 | 0.09 | 0.17 |
| Observations | 39,754 | 29,311 | 151,132 | 81,250 |

Ordinary Least Squares (OLS) is used to estimate each model which all include stack-by-player and stack-by-year fixed-effects. The depednent variable *Exit* is a binary estimate that is multiplied by one hundred for presentional simplicity. The mean of the dependent variable is calculated from pre-treatment observations of players in the treated group. *Clustered (Stack-by-Player) standard-errors in parentheses. Signif. Codes: ***: 0.01, **: 0.05, *: 0.1.*

among former teammates are all more pronounced in hockey than in football. These factors suggest that hockey players should have a more elastic labor supply response than football. Further, the relative time proximity difference in the timing of these deaths means that counter-factual groups can be more easily constructed in the hockey setting than for football. This can be seen in changes in the number of observations employed in the matched samples relative to the full data, falling 25~30% in the hockey setting and nearly 50% for football.

The coefficients of Table 3 reveal large treatment effect sizes. The preferred models, which use a matched control group (models two and four), reveal effect sizes ranging from 4.42-5.15 percentage points for hockey and 1.61-3.21 for football. These effect sizes correspond to 20-38% increases relative to the pre-treatment mean for hockey and 5-17% for football.

In order to investigate the dynamics of the effect of treatment across sports, I estimate two event study models via Equation 9 using the matched control group for hockey and football players. Estimates of pre-treatment differences in exit between the treated and control players are tightly centered around zero, providing evidence supporting the hypothesis that outcomes betwen the two groups would continue trending parallel to one another absent treatment. Following the timing of treatment, treated players in both sports see large, statistically significant increases in the probability of exit relative to the control group. However, both the magnitude and the dynamics of this post-treatment increase differ meaningfully across the sports, where the treatment effect for hockey is larger but also increases in magnitude over time. These results contrast with the smaller, transitorily treatment effect seen in football.

Notably, the treatment effect for both sports seems to be delayed one season, particularly in the case of hockey players whose estimates increase from approximately 1 percentage point in period zero to 4 percentage points in the following period. One potential reason for the (relatively) delayed labor supply response is the difference between the timing of a former teammate's death and their eventual CTE diagnosis.[25] Additionally, hockey contracts are longer on average than those seen in football. Thus, we would expect to see more delayed exit responses in hockey if individuals leave the sport after their contract at the time of treatment. I explore heterogeneity in contract structure as a driving force behind differences in retirement in the following subsection of this paper.

To further explore the mechanisms driving the heterogeneous labor supply responses that individuals display in response to the death of a coworker, I examine familiarity as an explanatory

[25]For instance, there was a delay of over 200 days between the date of death of Bob Probert (the first hockey player with teammates still in the league to be diagnosed with CTE) and his eventual diagnosis. Though delays of this magnitude are relatively common, speculation that the athlete died with CTE is more common in football, where the salience of the disease is more prevalent. For a more detailed discussion of this topic, please refer to Table A3 and its corresponding section in the Appendix.

**Figure 5.** Event Study of the Effect of the CTE Death of a Former Teammate on Professional Exit

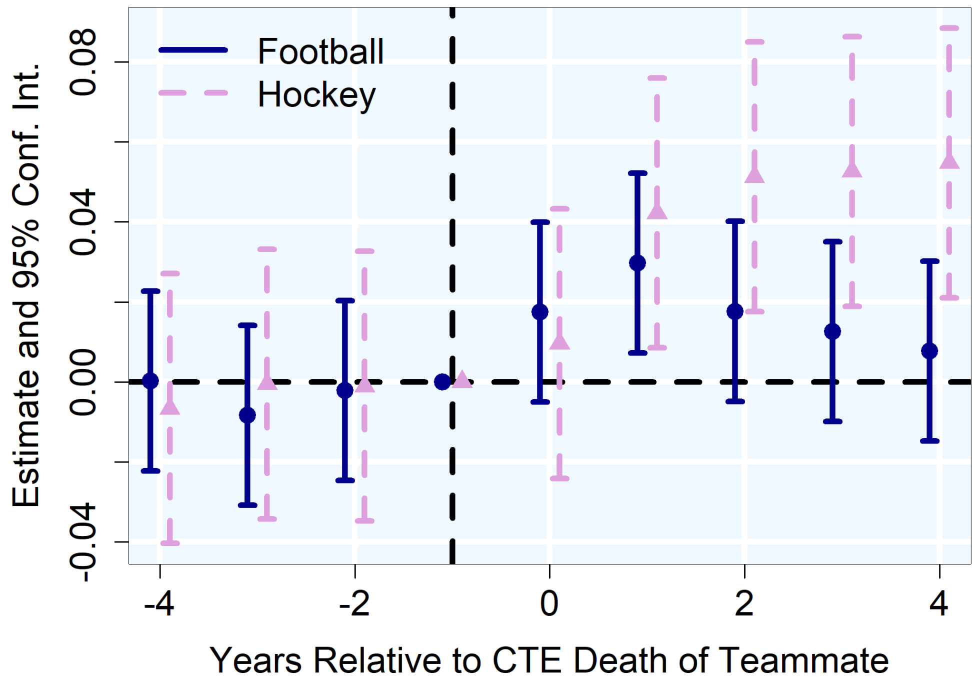


mechanism. These results can be seen in Table 4, which indicate that individual labor supply responses to peer health shocks exhibit a more elastic response as the amount of times individuals spent together on a team increases. In panel A of this table, the results are stratified by length of time spent as teammates of periods of one, two, and three years.[26] The results increase across each specification from 2.53 percentage points in model one to 4.67 percentage points in model three. This increase is even larger in percentage terms, rising from 9%-19% across the models as individuals who were teammates for longer have marginally lower probabilities of having retired prior to treatment.

As seen at the bottom of panel A of Table 4, individuals treated for longer periods have lower average pre-treatment differences in the share of those who have already exited and are slightly older. These age differences raise the possibility that the observed effect attributed to familiarity may instead reflect reversion to the mean, as older players are generally more likely to retire, all else being equal. Thus, I explore the role of heterogeneity in these findings across the age distribution in panel B. The models from this panel stratify the results by the age percentile that

[26]Over 60% of treated units were only teammates with a coworker who died of CTE for one season, as shown in Figure A8 in the Appendix. This magnitude of this relationship can be explained by a combination of high turnover in these sports and the young age at which their former teammates die. Over ninety percent of all treated players were on the same team for four or fewer seasons, which sharply limits the ability to stratify results beyond this cutoff without running into issues of statistical power.

**Table 4.** Peer Work-Related Deaths & Retirement: Heterogeneity by Teammate Duration

| Dependent Variable: | $\mathbb{1}$(Exit)*100 | | |
|---|---|---|---|
| Model: | (1) | (2) | (3) |
| **Panel A (Length of Time as Teammates)** | | | |
| CTE (Post) ×Teammate | 2.53*** | 4.30*** | 4.67*** |
| | (0.80) | (1.22) | (1.69) |
| Pre-Treatment Mean | 29.03 | 25.69 | 24.87 |
| % Change | 0.09 | 0.17 | 0.19 |
| Observations | 229,833 | 218,853 | 215,226 |
| Mean Pre-Treatment Age | 27.75 | 28.59 | 29.31 |
| Length of Time Teammates ($L$): | $L \geq 1$ | $L \geq 2$ | $L \geq 3$ |
| **Panel B (Teammates by Age)** | | | |
| CTE (Post) ×Teammate | 0.10 | 3.70** | 0.56 |
| | (2.22) | (1.53) | (1.72) |
| Pre-Treatment Mean | 10.8 | 17.31 | 36.27 |
| % Change | 0.01 | 0.21 | 0.02 |
| Observations | 97,794 | 131,022 | 32,643 |
| Pre-Treatment Mean Age | 22.36 | 26.03 | 32.25 |
| Age Percentile ($P$): | $P < 0.\bar{3}$ | $0.\bar{3} \leq P < 0.\bar{6}$ | $P \geq 0.\bar{6}$ |

Ordinary Least Squares (OLS) is used to estimate each model which all include stack-by-player and stack-by-year fixed-effects. Panel A subsets the treated group by the amount of time spent as teamamtes whereas panel B subsets the treated group by age quantiles. The age ranges for each model are 19-27, 28-31 and 32-48 years old. The depednent variable *Exit* is a binary estimate that is multiplied by one hundred for presentional simplicity. The mean of the dependent variable is calculated from pre-treatment observations of players in the treated group. *Clustered (Stack-by-Player) standard-errors in parentheses. Signif. Codes: ***: 0.01, **: 0.05, *: 0.1.*

players were in at the time of treatment consisting of individuals aged 18-25 (age percentile $(P) < \frac{1}{3}$), 26-29 ($\frac{1}{3} \leq P < \frac{2}{3}$) and 30 or older ($P \geq \frac{2}{3}$). Unlike in panel A of Table 4 where differences in the pre-treatment mean are negligible across each model, models one and three of panel A have differences in the pre-treatment mean of over twenty-five percentage points. Thus, in order to construct suitable counterfactual groups for each of these newly defined treated groups, I use the same iterative regression-based covariate balancing method described in the Methods section.[27]

The results of panel B of Table 4 reveal that, within this setting, the effect of workplace safety-related peer health shocks on extensive margin labor supply elasticities does not linearly increase with age. The estimates from models one and three, which analyze the youngest and oldest thirty-three percentiles of athletes who were exposed to the CTE death of a former teammate have statistically insignificant 0.1 (1%) and 0.56 (2%) percentage point increases in the probability of exiting the profession after their teammate's death relative to characteristically similar athletes. These effects contrast with the large and statistically significant 3.7 (21%) percentage point increase in the rate of exit among those in the middle of the age distribution. These findings reinforce the plausibility that the findings presented in panel A of this table are attributable to familiarity and not age.

Many factors could explain the findings from panel B of Table 4. For instance, differences in earnings potential between younger and mid-career athletes may play a significant role. Younger athletes, who are often still on their first ("rookie") contracts and whose wages are artificially constrained by binding price ceilings, face high expected opportunity costs of leaving the profession given the asymmetrically large salary benefits which can be earned after entering "free agency" where players can face more competition for their services. Mid-career athletes who have likely already secured contracts in free agency may face lower future expected opportunity costs and thus may be more sensitive to peer health shocks. In contrast, older athletes who have played for long periods may exhibit fatalistic risk preferences regarding CTE. Since CTE develops due to *cumulative* blows to the head, these athletes might feel that their risk is already cemented, regardless of whether they exit the profession.

Another plausible mechanism driving heterogeneity in the labor supply responses between sports may arise due to differences in the temporal proximity of the CTE death of a colleague. As can be seen in Figure 1, the discovery of CTE within hockey not only occurred far more suddenly than in football but also many of these individuals died much closer to the ends of their careers.[28]

[27] The event study estimates from these models can be seen in Figure A11 in the Appendix, providing evidence in favor of the hypothesis that the parallel pre-trends assumption required the estimation of treatment effects in difference-in-differences models has been satisfied.

[28] In the case of the three NHL players who died in 2011 and were later diagnosed with CTE, none had yet retired and each was on an active-roster.

To examine this hypothesis, I stratify the treated group based on the number of years between the last time an individual was on the same team with someone who died (and is later diagnosed with CTE) and the date of their death. This stratification allows one to test whether there is a gradient of labor supply responsiveness with respect to the temporal proximity of peer health shocks. Results from Table 5 provide evidence in favor of the hypothesis that individuals respond to peer health shocks more strongly when they are teammates more recently.

**Table 5.** Peer Work-Related Deaths & Retirement: Heterogeneity by Time since Last Treated

| Dependent Variable: | $\mathbb{1}$(Exit)*100 | | |
|---|---|---|---|
| Time since Treated (D): | $D \leq 2$ | $D \leq 4$ | $D \leq 6$ |
| Model: | (1) | (2) | (3) |
| **Panel A (Time since Treated)** | | | |
| CTE (Post) × Teammate | 3.14** | 1.78* | 1.91** |
| | (1.45) | (0.96) | (0.88) |
| Pre-Treatment Mean | 10.45 | 22.27 | 28.29 |
| % Change | 0.3 | 0.08 | 0.07 |
| Observations | 95,049 | 202,824 | 219,366 |
| **Panel B (Temp. Treat. Dist. & 3+ Years Teammates)** | | | |
| CTE (Post) × Teammate | 7.22* | 4.62* | 3.81* |
| | (4.05) | (2.65) | (2.29) |
| Pre-Treatment Mean | 10.45 | 22.27 | 28.29 |
| % Change | 0.69 | 0.21 | 0.13 |
| Observations | 39,231 | 68,256 | 73,035 |

Ordinary Least Squares (OLS) is used to estimate each model, which includes stack-by-player and stack-by-year fixed-effects. The dependent variable *Exit* is a binary estimate multiplied by one hundred for presentational simplicity. The mean of the dependent variable is calculated from pre-treatment observations of players in the treated group. *Clustered (Stack-by-Player) standard errors in parentheses. Significance Codes: ***: 0.01, **: 0.05, *: 0.1.*

Table 5 examines heterogeneous peer effects of the impact of the CTE death of a former teammate on the probability of exiting the profession by the length of time between when individuals last played together for a team and the date of their eventual death. These treated individuals are stratified into periods of two, four, and six years, where each subgroup is contained in the larger year bin. Relative to a group of characteristically similar athletes prior to treatment, treated athletes who were last teammates with someone who died two or fewer years ago saw a 3.14 percentage point increase in the likelihood of exiting. Relative to a lower pre-treatment retirement rate, this represents a large 30% % increase. Model two (three) shows that this effect size decays when including individuals that shared a team with at least four (six) years prior to their teammate's death to 1.78 (1.91) percentage points. Relative to the pre-treatment mean, this estimate reflects an 8% (7%) increase in the probability of retiring.

Panel B of table 5 jointly examines the interactions of temporal proximity and familiarity with peers on the responsiveness of their teammates' labor supply. To test this interaction, the treated groups from panel A are further stratified by individuals treated for at least two years across each model in this panel. The results reveal interactions that are greater than the sum of their parts. For instance, the estimates of model one show that athletes who were teammates with a former teammate who died with CTE for three or more years and played for a team with them at least two years before their death are 7.22 percentage points more likely to retire than characteristically similar non-treated players in the same years. Relative to the pre-treatment mean, this represents a 69% increase. This effect size is meaningfully larger than the sum of the 30% effect of proximity of treatment seen in model one of panel A in Table 5 and the 19% effect of familiarity seen in model three of panel A in Table 4. The addition of individuals treated four (six) years ago or fewer reduces the effect sizes to 4.62 (3.81) percentage points which represent a 21% (13%) increases relative to the pre-treatment mean.

To test whether the treatment effects reported thus far can more plausibly be attributed to factors such as the increased salience of risk or any other range of emotional trials which could be experienced following the death of a coworker, I perform a placebo analysis where CTE deaths are replaced with deaths that are unrelated to workplace safety. I chose the two most common causes of death for young athletes: car accidents and rare-genetic (typically cardiac-related) diseases.[29] The effects of this exercise can be seen in Table 6. The results from this table fail to reject the hypothesis that deaths unrelated to workplace safety differentially increase former teammates' likelihood of retiring.These findings bolster the argument that peer health shocks increase retirements through the channel of increases in the salience of risk.

### 8.b. Contracts

I next turn my attention to studying the subsequent employment contracts of treated athletes who remain playing risky professional sports after the CTE death of a former teammate. Economic theory posits that if individuals perceive types of work as relatively less desirable, then workers will exit this profession or workplace until wages rise to compensate for the undesirability of the job. Alternatively, this theory suggests that workers in more hazardous jobs who are willing to continue working do so because they are able to receive greater compensation which help to offset the utility loss from have to face the workplace risks which have increased in salience.

[29]Note to Josh: I need a reference to an appendix table here with these names.

**Table 6.** Peer Deaths Unrelated to Workplace Safety & Retirement: A Placebo Test

| Dependent Variable: | 𝟙(Exit)*100 | |
|---|---|---|
| Model: | (1) | (2) |
| *Variables* | | |
| Car Accident (Post) × Teammate | -0.01 | |
| | (0.01) | |
| Disease (Post) × Teammate | | 0.00 |
| | | (0.01) |
| Mean of Dep. | 0.19 | 0.19 |
| % Change | -0.06 | 0.00 |
| Observations | 95,533 | 70,891 |

Each model includes Stack-by-Player and Stack-by-Year fixed-effects. *Clustered (Stack-by-Player) standard-errors in parentheses. Signif. Codes: ***: 0.01, **: 0.05, *: 0.1.*

In a market where wages can freely and rapidly adjust to the changes in preferences of its workers, athletes' voluntary exits from the profession should occur only in settings where the marginal value of the labor of the worker (and thus in the wages offered from the firm) fall bellow the worker's reservation wage.[30] However, in this setting, worker compensation is very heavily regulated.[31] These binding price ceilings mute variation in worker compensation relative to what would be expected in less regulated markets, which, in turn, increases the rate at which workers exit the profession.

To study the impact of peer-health-shock-induced increases in risk perception on the details of workers' subsequent employment contracts, I estimate a series of difference-in-differences models in Table 7. These models most notably differ from those in the previous subsection in their use of dependent variables. These new measures include total compensation, guaranteed compensation, salary, and the length of time individuals are contracted with a team. Another important change in these models is the switch from the usage of balanced data to unbalanced data, which can impact the interpretation of some coefficients if these measures capture realized compensation rather than what is enumerated in the contract. These measures of compensation will equal one another if a player is not released by the team before the end of their contract. Finally, I subset this data to only include observations in the contracts that individuals sign immediately before and (if applicable) after treatment. Reducing the sample size by this amount helps to isolate treatment

[30] An involuntary exit refers to situations in which workers are cut from team rosters when the marginal cost of their employment is greater than the marginal benefit.

[31] See the Setting section (Section 3) of this paper for more information on how wage controls differ between the NFL and NHL.

effects more accurately by reducing variation from contracts that are signed much earlier or later in one's career.

Each model in 7 is estimated using Poisson regression, whose coefficients are interpreted as the log of the expected count ratio. These coefficients can be closely approximated for small changes as percentage changes in the outcome. Furthermore, I condition the treatment effects in models one, two, and three on the length of the contract. This approach ensures that the treatment effects are more directly attributable to changes in the outcome variable itself rather than being influenced by variations in contract length, as would happen if the estimates were annualized by dividing by the contract duration. The results from this table demonstrate that conditional upon not exiting, treated athletes exchange reductions in annual salary for increased signing bonuses and shorter contracts. On the net, this moderately increases total compensation, albeit statistically insignificant.

**Table 7.** The Effect of the CTE Death of a Former Teammate on Subsequent Contract Details

| Dependent Variables:<br>Model: | Total<br>(1) | Guaranteed<br>(2) | Salary<br>(3) | Length<br>(4) |
|---|---|---|---|---|
| **Panel A (Full Sample)** | | | | |
| CTE (Post) ×Teammate ×New Ct. | 0.07 | $0.29^{***}$ | $-0.28^{***}$ | $-0.07^{*}$ |
| | (0.07) | (0.07) | (0.06) | (0.04) |
| Length | $0.37^{***}$ | $0.23^{***}$ | $0.13^{***}$ | |
| | (0.01) | (0.01) | (0.01) | |
| Pre-Treatment Mean | 13.84 | 1.17 | 1.76 | 3.84 |
| Change | 1.03 | 0.34 | -0.49 | -0.29 |
| Observations | 77,987 | 61,207 | 77,945 | 78,652 |
| **Panel B (Balanced Sample)** | | | | |
| CTE (Post) ×Teammate ×New Ct. | 0.02 | $0.22^{***}$ | $-0.28^{***}$ | $-0.07^{*}$ |
| | (0.07) | (0.07) | (0.06) | (0.04) |
| Length | $0.37^{***}$ | $0.23^{***}$ | $0.13^{***}$ | |
| | (0.01) | (0.01) | (0.01) | |
| Pre-Treatment Mean | 14.38 | 1.21 | 1.86 | 4.03 |
| Change | 0.28 | 0.26 | -0.52 | -0.29 |
| Observations | 52,912 | 44,595 | 52,909 | 53,271 |

Poisson regression is used to estimate each model. Each model includes Stack-by-Player and Stack-by-Year fixed-effects. *Clustered (Stack-by-Player) standard-errors in parentheses.* *Signif. Codes: ***: 0.01, **: 0.05, *: 0.1.*

Model one in panel A of Table 7 measures the degree to which total compensation changes in the treated group relative to the control group before and after treatment. This variable reflects that which is enumerated in the contract, and thus, treatment effect estimates of this variable are not influenced by the amount of compensation that the player ultimately receives. The coefficient reveals that treated players' total compensation rose approximately 7% relative to a control group after treatment. This represents an increase of approximately $1.03 million. While $1 million represents a non-trivial increase in compensation, this estimate is imprecisely estimated and masks more interesting variations in the specifics of treated players' contracts.

Model two in panel A of Table 7 measures the amount of additional guaranteed compensation that players receive due to the increased salience of workplace safety. The coefficient reveals that treated players' signing bonuses increased significantly by approximately 30 percent. This substantial increase in signing bonuses, relative to the pre-treatment mean, represents an increase in total compensation of $340k. This boost in financial security is a direct result of the increased focus on workplace safety. Similar to total compensation in model one, guaranteed compensation is owed to workers regardless of whether they are still employed by the team in the future. Thus, the estimates from this model are not conditional upon the continued employment of the worker beyond any point after which they sign the contract. This contrasts with estimates from model three of this table which measures how total salary changes due to treatment. The coefficient of interest from this model reveals that the total amount of salary earned by players in their contracts signed after the CTE death of a former teammate is approximately 30% lower than expected. In order for the findings from models one, two and three to be simultaneously true, it must be the case that treated players are disproportionately likely to exit the profession before the end of the contract signed after treatment.

Finally, model four in panel A of Table 7 reveals that treated workers sign contracts that are approximately 7% shorter than individuals in the control group after treatment. Relative to the pre-treatment mean, this corresponds to a reduction in contract length of approximately one-third of a year. The results from these models in panel A present a clear story of risk mitigation strategies on the part of athletes with a greater salience of risk after the CTE death of a former teammate – they sign shorter contracts, receive large increases in their signing bonus, and then exit the profession earlier than individuals in the control group while forgoing less salary for doing so.

In order to test whether the results from panel A of Table 7 are driven by endogenous selection of workers who remain in the profession, I restrict both the treated and control group only to include workers who sign contracts at least once before and after treatment. Though this significantly reduces the sample sizes available for estimation, the magnitude and statistical significance of the estimated treatment effects remain remarkably consistent across panels. The only estimated effect that meaningfully reduces is that for total compensation, which falls from 7% to 2% while remaining statistically insignificant.

As is commonly known, the major identifying assumption of difference-in-differences models is that, absent treatment, the outcomes of treated and control groups would continue to trending parallel to one another. Thus, I test for the plausibility of this assumption via an event study analysis of two contract-related outcome variables – salary and a binary indicator for whether individuals have any amount of guaranteed compensation within their contract. Limiting observations to individual's employment contracts before and, if applicable, after treatment, Figure 6 that the outcomes between treated and control players were trending parallel to one another prior to treatment. These outcomes diverge after treatment with the probability of treated players having any amount of their compensation guaranteed increasing by approximately 5% and their total earned salary decreasing approximately 10%.

**Figure 6.** Event Study of the Effect of the CTE Death of a Former Teammate on Subsequent Contract Details

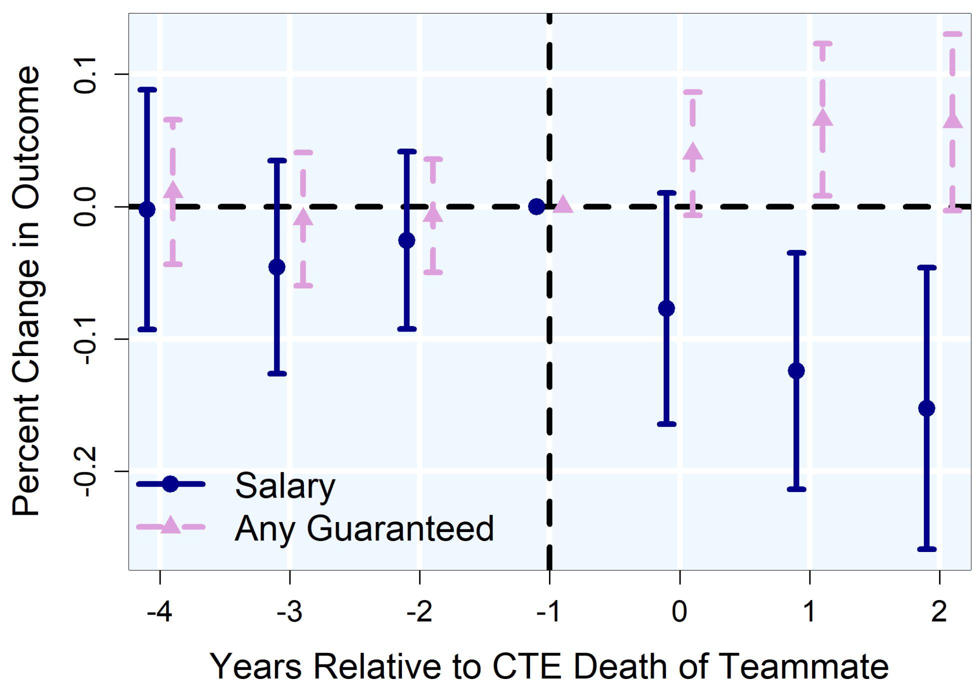


I next separately estimate contract-related treatment effects by sport in Table 8. Panel A of this table estimates effects for hockey players using the same models from Table 7. Each treatment coefficient in this panel is both statistically insignificant and less than ten percent in absolute magnitude, suggesting that treated individuals were unable to meaningfully alter the terms of their subsequent employment contracts in order to satisfactorily internalize the updated salience of the risks of continued workplace participation.

The null results from panel A of 8 are potentially unsurprising for two reasons. Most notably,

**Table 8.** The Effect of the CTE Death of a Former Teammate on Subsequent Contract Details (by Sport)

| Dependent Variables: | Total | Guaranteed | Salary | Length |
|---|---|---|---|---|
| Model: | (1) | (2) | (3) | (4) |
| **Panel A (Hockey)** | | | | |
| CTE (Post) ×Teammate ×New Ct. | -0.01 | -0.08 | -0.04 | 0.08 |
| | (0.09) | (0.42) | (0.06) | (0.09) |
| Length | 0.28*** | 0.40*** | 0.10*** | |
| | (0.02) | (0.07) | (0.01) | |
| Pre-Treatment Mean | 16.84 | 0.15 | 2.98 | 4.23 |
| Change | -0.21 | -0.01 | -0.12 | 0.33 |
| Observations | 10,980 | 4,093 | 10,972 | 11,278 |
| **Panel B (Football)** | | | | |
| CTE (Post) ×Teammate ×New Ct. | 0.13 | 0.29*** | -0.38*** | -0.14*** |
| | (0.08) | (0.07) | (0.08) | (0.04) |
| Length | 0.39*** | 0.22*** | 0.14*** | |
| | (0.01) | (0.01) | (0.01) | |
| Pre-Treatment Mean | 13.42 | 1.32 | 1.59 | 3.78 |
| Change | 1.8 | 0.38 | -0.61 | -0.55 |
| Observations | 67,007 | 57,114 | 66,973 | 67,374 |

Poisson regression is used to estimate each model. Each model includes Stack-by-Player and Stack-by-Year fixed-effects. *Clustered (Stack-by-Player) standard-errors in parentheses. Signif. Codes: ***: 0.01, **: 0.05, *: 0.1.*

in exchange for more player-friendly concessions such as longer, mostly-guaranteed contracts relative to other professional sports leagues, the salary cap is far more binding in hockey.[32] This cap strongly limits player flexibility in contract negotiations due to these caps which restrict annual total team wage expenditures. Second, though only six players on the ice at any one time per team in hockey, the high frequency with which hockey players rotate throughout the game means that there is less ability for any one individual to determine the outcome matches relative to other sports. This limits superstar effects within the sport which, in turn, makes the terms of contracts more uniform.[33]

The lack of significant results from panel A of 8 contrast strongly with those of panel B which analyze treatment effects for football players. Consistent with previous findings, model one of this panel estimates statistically insignificant effects of treatment on total compensation. However, the magnitude is meaningfully different with an estimate of a 13% increase. Relative to the pre-treatment mean, this represents a \$1.8 million increase. Similar to what was seen in table 7, salaries and contract lengths decrease approximately 40% and 14% (\$610 thousand and 0.55 years) while guaranteed compensation increases 30% (\$380k).

Given that the binding nature of salary caps reduces variation in total compensation needed to estimate compensating wage differentials in model one of Tables 8 and 7, I turn my attention to alternative identification strategy which leverages quasi-random variation in the amount of salary owed to athletes at the time of the CTE death of a teammate.

An analysis of the impact of opportunity costs on player retirement decisions can be seen in Table 9. Model one and two of this table estimate these effects on the full sample, finding that significant impacts of both treatment, opportunity costs and their interaction. Specifically, estimates from model one reveal that, holding all else equal, treatment increases the probability of exit by approximately 3.6 percentage points while a one percent increase in the amount of money individuals would forgo upon retirement decreases this probability by approximately 0.5 percentage points. The coefficient of interest in this table which interacts these two variables reveals treated workers' labor supply decisions are highly sensitive to opportunity costs. The coefficient for this interaction reveals that treated workers see an additional 0.37 percentage increase in the probability of retiring for every one percentage point increase in their opportunity cost.

Model two of Table 9 extends the findings of model one by adding two key control variables: the total amount of compensation agreed to on the contract and the number of years it covers.

[32] See section 3 (Setting) for a more lengthy discussion of the NHL salary cap.

[33] See Figure A1e for a visualization of non-standardized differences productivity in across hockey and football. The modal player in the hockey is far more likely to have no impact on the game, while the distribution of football player's productivity has a very long right-tail, suggesting a greater number of superstars.

**Table 9.** Peer Work-Related Deaths & Retirement: Heterogeneity by Opportunity Cost

| Dependent Variable: | $\mathbb{1}$(Exit)*100 | | | |
|---|---|---|---|---|
| Sample: | Full | | Hockey | Football |
| Model: | (1) | (2) | (3) | (4) |
| *Variables* | | | | |
| CTE (Post) ×Teammate | 3.59** | 2.48* | 1.53 | 4.08** |
| | (1.45) | (1.46) | (1.41) | (1.88) |
| $sinh^{-1}$(Opp. Cost) | -0.48*** | -0.06*** | -0.09*** | -0.02 |
| | (0.01) | (0.01) | (0.01) | (0.02) |
| CTE (Post) ×Teammate × $sinh^{-1}$(Opp. Cost) | -0.37*** | -0.32*** | -0.06 | -0.52*** |
| | (0.09) | (0.09) | (0.08) | (0.12) |
| $sinh^{-1}$(Total) | | -1.93*** | 0.95*** | -2.92*** |
| | | (0.13) | (0.16) | (0.18) |
| Length | | -2.83*** | -0.66*** | -3.06*** |
| | | (0.10) | (0.08) | (0.13) |
| Observations | 131,424 | 131,424 | 18,679 | 112,745 |

Ordinary Least Squares (OLS) is used to estimate each model. Each model includes Stack-by-Player and Stack-by-Year fixed-effects. *Clustered (Stack-by-Player) standard-errors in parentheses. Signif. Codes: ***: 0.01, **: 0.05, *: 0.1.*

These control variables help to isolate the effect of the opportunity cost on retirement by comparing individuals with similar contract details but only differ in the amount of money remaining on their contract at the time of treatment. While adding these control variables unsurprisingly reduces the magnitude of the coefficient for opportunity cost, it does little to change the effect size of its interaction of treatment.

Models three and four of Table 9 extend the estimating equation used in model two to separately present the differential impacts of opportunity costs on retirement for athletes in hockey and football. Model three finds little evidence that treated players with larger amounts of money remaining on their contract at the time of treatment in hockey were differentially more likely to exit after the CTE death of a former teammate. This is consistent with the knowledge that in hockey, players' remaining yet-paid salaries are far more likely to be received following their retirement due to the terms of their collective bargaining agreement. This contrasts sharply with football, whose athletes are far less likely to receive their yet-paid salaries upon retiring. Estimates from model four reveal that a one percent increase in treated football players' opportunity cost is estimated to reduce their probability of exiting by 0.5 percentage points relative to non-treated players with similar contracts.

To conclude, I utilize the predicted values from model two of Table 9 to estimate the dollar amount that treated individuals would have to receive to be indifferent between retirement and remaining within the workforce. This regression model provides me with a continuous measure of the predicted probability of exit from the profession based on the length of their contract, the total amount of compensation specified within it, their opportunity cost and their treatment status. I present the results from this exercise in Figure 7.

**Figure 7.** Opportunity Costs and the Probability of Exit

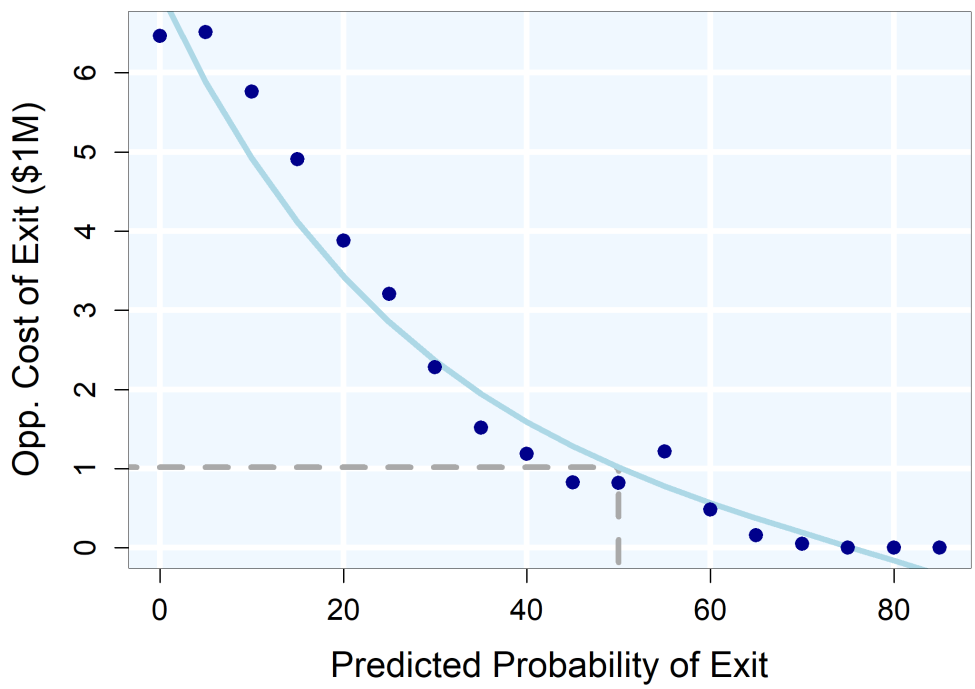


Figure 7 displays the relationship between opportunity cost and the predicted exit probability. The estimates are aggregated into bins, five percentage points in width for visual simplicity. The indifference point on the figure can be seen at the intersection of the line of best fit for these points (seen as a downward-sloping light blue shade in the figure), and the predicted probability of exit value is equal to 50. This intersects the y-axis at a value approximately equal to $1 million.[34] However, reading left-to-right across this entire figure reveals that, on average, athletes would need to be compensated an additional $6 million annually if teams wanted to know the minimum amount that they would have to pay and be very confident in preventing their workers from exiting.

While the downward-sloping nature of this curve matches theoretical priors, its convexity is potentially surprising. If individuals receive diminishing marginal utility from income, then one would expect that the estimates from Figure 7 would be concave. The intuition behind this reasoning is that moving from zero additional dollars of income to one million additional dollars

[34]Interestingly, this estimate almost perfectly equals the compensating wage differential estimate from model one in panel A of Table 7.

should increase one's probability of remaining in the profession by a greater amount than moving from one million dollars of income to two million.

On the other hand, the convexity of the curve could be a mechanical artifact of the y-axis since opportunity costs of exit cannot fall below zero, causing the curve to steepen as it approaches this lower bound. Additionally, athletes may perceive risks in a nonlinear fashion, requiring significantly more compensation as the perceived health risks of remaining in the profession grow after a peer-health shock. This heightened risk perception could lead to an exponential increase in the required compensation to deter exit, explaining the convex shape of the relationship.

## 9 Conclusion

In examining the effects of informational health shocks on labor supply, I find that former peers exhibit stronger responses than expected, suggesting that individuals place greater weight on information relayed through their social networks, even in settings where risk is highly publicized. This finding contradicts the predictions of standard economic models, which typically assume that individuals fully incorporate all available knowledge.

This paper draws upon and contributes to insights from behavioral economics, particularly concerning biases related to personal health status and associated risks. Research shows that individuals commonly exhibit optimistic biases regarding personal risk (Weinstein, 1989). Such biases persist even in settings where information is easily accessible and errors can have serious financial and health consequences (Bhattacharya et al., 2009; Golman et al., 2017; Oster et al., 2013). The prevalence of these biases has important implications for health and wage inequalities among workers, as well as for policies concerning optimal provision of workers' compensation (Viscusi, 1980). Further, these findings contribute to existing work in behavioral economics by aligning with evidence that geographic proximity to events, such as natural disasters or abnormal weather patterns, significantly impacts risk perceptions (Egan & Mullin, 2012; McCoy & Walsh, 2018). Similarly, the literature on gender and racial concordance in health economics indicates that individuals are more receptive to advice from medical professionals who share their racial or gender background (Alsan et al., 2019, 2024; Alsan & Eichmeyer, 2024; Cabral & Dillender, 2024). Additionally, the pronounced effect of treatment recency aligns with findings from meta-analyses showing that discounts on flood-risk housing are highly sensitive to the time elapsed since the last flood (Beltrán et al., 2018). Together, these insights underscore how familiarity, shared experience, and recency can intensify responses to risk – a peer's health outcome may hold unique salience due to the close connection and demographic similarity shared by teammates.

There are numerous reasons to believe that the estimates presented in this paper represent a lower bound of the "true" effect of peer influences on risk perception. First, athletic talent is often observable early in one's life, leading to substantial investments in sport-specific skills from a young age. These skills are not easily transferable to other fields, thus limiting athletes' ability to credibly consider voluntary exit from their profession. Second, the strong effects observed here suggest that peer health shocks may have an even greater influence on labor supply decisions than previously assumed, given that CTE and concussion risks are widely covered in the media and affect all players. Third, the young, competitive, and entirely male sample of workers analyzed in this study likely have significantly higher baseline risk tolerance than the general population which may reduce responsiveness to changes in risk salience (Croson & Gneezy, 2009; Schildberg-Hörisch, 2018).

While this setting offers unique insights into the mechanisms behind peer effects, it is limited in its ability to assess whether workers are "appropriately" responding to risk. Existing literature has shown that individuals tend to over-extrapolate from personal experiences and that changes in risk perception may be short-lived (Kuchler & Zafar, 2019; McCoy & Walsh, 2018). However, gauging rationality in this context would require reliable estimates of the actual risk of developing CTE, which remain elusive. The inability to diagnose CTE in living individuals, along with significant endogeneity in posthumous testing, has led to wide variation in estimated risk. For instance, Boston University's Chronic Traumatic Encephalopathy Center found that over 90% of the nearly 400 former athletes studied posthumously showed signs of CTE while a recent survey of over 4,000 former professional football players found that "only" one-third believe they are currently living with the disease (Chobanian & Avedisian School of Medicine, 2023; Grashow et al., 2024). However, given the substantial financial stakes involved in retirement decisions for professional athletes, it seems more plausible that these workers would update their beliefs toward the "true" mortality risk rather than abruptly overestimating it.

To conclude, this study advances the literature on peer effects and labor economics by illustrating the limited ability of wage adjustments alone to retain workers in high-risk environments, where individuals tend to underestimate the utility loss associated with work-related health risks. This underestimation often leads to suboptimal career choices and increased exposure to workplace hazards that are not adequately compensated, ultimately putting individual well-being and workforce stability at risk. Over time, these misjudgments may also reduce aggregate productivity by impacting workforce efficiency and retention. These findings highlight the importance of more effective communication about occupational health risks, especially through trusted messengers who have close ties to workers. Policy implications include the potential for targeted awareness campaigns and safety-oriented provisions in collective bargaining agreements, which could better align workers' risk perceptions with actual hazards. Future research could explore how variations in contract structures across industries shape workers' decisions to remain in high-risk roles, and assess how evolving safety measures influence the balance between compensation and risk perception in labor markets.

## References


Adgate, B. (2024). *NFL games made up 93% of the most watched TV programs in 2023*. https://www.forbes.com/sites/bradadgate/2024/01/05/nfl-games-made-up-93-of-the-most-watched-tv-programs-in-2023/

Adler, M. (1985). Stardom and talent. *The American Economic Review*, *75*(1), 208–212.

Alsan, M., Durvasula, M., Gupta, H., Schwartzstein, J., & Williams, H. (2024). Representation and extrapolation: Evidence from clinical trials. *The Quarterly Journal of Economics*, *139*(1), 575–635.

Alsan, M., & Eichmeyer, S. (2024). Experimental evidence on the effectiveness of nonexperts for improving vaccine demand. *American Economic Journal: Economic Policy*, *16*(1), 394–414.

Alsan, M., Garrick, O., & Graziani, G. (2019). Does diversity matter for health? Experimental evidence from oakland. *American Economic Review*, *109*(12), 4071–4111.

Anderson, K. H., & Burkhauser, R. V. (1985). The retirement-health nexus: A new measure of an old puzzle. *Journal of Human Resources*, 315–330.

Anderson, P. (2022). Compensating differentials for the risk of reinjury–lessons from professional boxing. *Journal of Sports Economics*, 15270025221120592.

Bailey, M., Cao, R., Kuchler, T., & Stroebel, J. (2018). The economic effects of social networks: Evidence from the housing market. *Journal of Political Economy*, *126*(6), 2224–2276.

Beltrán, A., Maddison, D., & Elliott, R. J. (2018). Is flood risk capitalised into property values? *Ecological Economics*, *146*, 668–685.

Bender, K. A., & Mridha, H. (2011). The effect of local area unemployment on compensating wage differentials for injury risk. *Southern Economic Journal*, *78*(2), 287–307.

Bhattacharya, J., Goldman, D., & Sood, N. (2009). Market evidence of misperceived mortality risk. *Journal of Economic Behavior & Organization*, *72*(1), 451–462.

Black, D. (1986). *The plague years: A chronicle of AIDS, the epidemic of our times*. Simon; Schuster New York.

Bleemer, Z., & Zafar, B. (2018). Intended college attendance: Evidence from an experiment on college returns and costs. *Journal of Public Economics*, *157*, 184–211.

Böckerman, P., & Ilmakunnas, P. (2009). Job disamenities, job satisfaction, quit intentions, and actual separations: Putting the pieces together. *Industrial Relations: A Journal of Economy and Society*, *48*(1), 73–96.

Böckerman, P., & Ilmakunnas, P. (2020). Do good working conditions make you work longer? Analyzing retirement decisions using linked survey and register data. *The Journal of the Economics of Ageing*, *17*, 100192.

Brown, J. R. (2001). Private pensions, mortality risk, and the decision to annuitize. *Journal of*

*Public Economics*, *82*(1), 29–62.

Brown, J. R., Coile, C. C., & Weisbenner, S. J. (2010). The effect of inheritance receipt on retirement. *The Review of Economics and Statistics*, *92*(2), 425–434.

Cabral, M., & Dillender, M. (2024). Gender differences in medical evaluations: Evidence from randomly assigned doctors. *American Economic Review*, *114*(2), 462–499.

Cardazzi, A., Martin, J. C., & Rodriguez, Z. (2023). Information shocks and celebrity exposure: The effect of "Magic" Johnson on AIDS diagnoses and mortality in the US. *Health Economics*, *32*(9), 2047–2079.

Carrell, S. E., Fullerton, R. L., & West, J. E. (2009). Does your cohort matter? Measuring peer effects in college achievement. *Journal of Labor Economics*, *27*(3), 439–464.

Carrell, S. E., Hoekstra, M., & Kuka, E. (2018). The long-run effects of disruptive peers. *American Economic Review*, *108*(11), 3377–3415.

Carrell, S. E., Hoekstra, M., & West, J. E. (2011). Is poor fitness contagious?: Evidence from randomly assigned friends. *Journal of Public Economics*, *95*(7-8), 657–663.

Carrell, S. E., Hoekstra, M., & West, J. E. (2019). The impact of college diversity on behavior toward minorities. *American Economic Journal: Economic Policy*, *11*(4), 159–182.

Chen, J., & Roth, J. (2024). Logs with zeros? Some problems and solutions. *The Quarterly Journal of Economics*, *139*(2), 891–936.

Chetty, R., Friedman, J. N., Hilger, N., Saez, E., Schanzenbach, D. W., & Yagan, D. (2011). How does your kindergarten classroom affect your earnings? Evidence from project STAR. *The Quarterly Journal of Economics*, *126*(4), 1593–1660.

Chobanian & Avedisian School of Medicine. (2023). *Researchers find CTE in 345 of 376 former NFL players studied*. https://www.bumc.bu.edu/camed/2023/02/06/researchers-find-cte-in-345-of-376-former-nfl-players-studied/

Chung, B. W., & Zou, J. (2023). Understanding spillover of peer parental education: Randomization evidence and mechanisms. *Economic Inquiry*, *61*(3), 496–522.

Cohen-Zada, D., Dayag, I., & Gershoni, N. (2024). Effort peer effects in team production: Evidence from professional football. *Management Science*, *70*(4), 2355–2381.

Coile, C. C., & Levine, P. B. (2007). Labor market shocks and retirement: Do government programs matter? *Journal of Public Economics*, *91*(10), 1902–1919.

Coile, C. C., & Levine, P. B. (2011). Recessions, retirement, and social security. *American Economic Review*, *101*(3), 23–28.

Cook, C., Diamond, R., Hall, J. V., List, J. A., & Oyer, P. (2021). The gender earnings gap in the gig economy: Evidence from over a million rideshare drivers. *The Review of Economic Studies*, *88*(5), 2210–2238.

Cornelissen, T., Dustmann, C., & Schönberg, U. (2017). Peer effects in the workplace. *American*

*Economic Review*, *107*(2), 425–456.

Cottini, E., Kato, T., & Westergaard-Nielsen, N. (2011). Adverse workplace conditions, high-involvement work practices and labor turnover: Evidence from danish linked employer–employee data. *Labour Economics*, *18*(6), 872–880.

Croson, R., & Gneezy, U. (2009). Gender differences in preferences. *Journal of Economic Literature*, *47*(2), 448–474.

Dessaint, O., & Matray, A. (2017). Do managers overreact to salient risks? Evidence from hurricane strikes. *Journal of Financial Economics*, *126*(1), 97–121.

Doepke, M., & Zilibotti, F. (2008). Occupational choice and the spirit of capitalism. *The Quarterly Journal of Economics*, *123*(2), 747–793.

Dohmen, T., Falk, A., Huffman, D., & Sunde, U. (2012). The intergenerational transmission of risk and trust attitudes. *The Review of Economic Studies*, *79*(2), 645–677.

Dole, C., & Kassis, M. M. (2010). Evidence of a compensating wage differential for NFL players who play on artificial turf. *International Journal of Sport Management and Marketing*, *8*(3-4), 254–264.

Duflo, E., & Saez, E. (2003). The role of information and social interactions in retirement plan decisions: Evidence from a randomized experiment. *The Quarterly Journal of Economics*, *118*(3), 815–842.

Egan, P. J., & Mullin, M. (2012). Turning personal experience into political attitudes: The effect of local weather on americans' perceptions about global warming. *The Journal of Politics*, *74*(3), 796–809.

Eisenberg, D., Golberstein, E., & Whitlock, J. L. (2014). Peer effects on risky behaviors: New evidence from college roommate assignments. *Journal of Health Economics*, *33*, 126–138.

Fainaru-Wada, M., & Fainaru, S. (2013). *League of denial: The NFL, concussions, and the battle for lruth.* Crown.

Falk, A., Becker, A., Dohmen, T., Enke, B., Huffman, D., & Sunde, U. (2018). Global evidence on economic preferences. *The Quarterly Journal of Economics*, *133*(4), 1645–1692.

Falk, A., & Ichino, A. (2006). Clean evidence on peer effects. *Journal of Labor Economics*, *24*(1), 39–57.

Feld, J., & Zölitz, U. (2017). Understanding peer effects: On the nature, estimation, and channels of peer effects. *Journal of Labor Economics*, *35*(2), 387–428.

Fishback, P. V. (1986). Workplace safety during the progressive era: Fatal accidents in bituminous coal mining, 1912–1923. *Explorations in Economic History*, *23*(3), 269–298.

Fishback, P. V., & Kantor, S. E. (1992). "Square deal" or raw deal? Market compensation for workplace disamenities, 1884–1903. *The Journal of Economic History*, *52*(4), 826–848.

Fishback, P. V., & Kantor, S. E. (1995). Did workers pay for the passage of workers' compensation laws? *The Quarterly Journal of Economics*, *110*(3), 713–742.

Galor, O., & Özak, Ö. (2016). The agricultural origins of time preference. *American Economic Review*, *106*(10), 3064–3103.

Giustinelli, P., & Shapiro, M. D. (2024). SeaTE: Subjective ex ante treatment effect of health on retirement. *American Economic Journal: Applied Economics*, *16*(2), 278–317.

Golman, R., Hagmann, D., & Loewenstein, G. (2017). Information avoidance. *Journal of Economic Literature*, *55*(1), 96–135.

Grashow, R., Terry, D. P., Iverson, G. L., DiGregorio, H., Dairi, I., Brown, C., Atkeson, P. S., Whittington, A. J., Reese, L., & Kim, J. H. (2024). Perceived chronic traumatic encephalopathy and suicidality in former professional football players. *JAMA Neurology*.

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

Hanaoka, C., Shigeoka, H., & Watanabe, Y. (2018). Do risk preferences change? Evidence from the great east japan earthquake. *American Economic Journal: Applied Economics*, *10*(2), 298–330.

Hanson, A., Jolly, N. A., & Peterson, J. (2017). Safety regulation in professional football: Empirical evidence of intended and unintended consequences. *Journal of Health Economics*, *53*, 87–99.

Hersch, J. (1998). Compensating differentials for gender-specific job injury risks. *The American Economic Review*, *88*(3), 598–607.

Hersch, J. (2011). Compensating differentials for sexual harassment. *American Economic Review*, *101*(3), 630–634.

Hoxby, C. M. (2000). *Peer effects in the classroom: Learning from gender and race variation* (Working Paper 7867). National Bureau of Economic Research.

Hu, Z. (2022). Social interactions and households' flood insurance decisions. *Journal of Financial Economics*, *144*(2), 414–432.

Imbens, G. W., Rubin, D. B., & Sacerdote, B. I. (2001). Estimating the effect of unearned income on labor earnings, savings, and consumption: Evidence from a survey of lottery players. *American Economic Review*, *91*(4), 778–794.

Innocenti, S., Clark, G. L., McGill, S., & Cuñado, J. (2019). The effect of past health events on intentions to purchase insurance: Evidence from 11 countries. *Journal of Economic Psychology*, *74*, 102204.

Kaul, A., Abbas, A., Smith, G., Manjila, S., Pace, J., & Steinmetz, M. (2016). A revolution in preventing fatal craniovertebral junction injuries: Lessons learned from the head and neck support device in professional auto racing. *Journal of Neurosurgery: Spine*, *25*(6), 756–761.

Kim, S.-W., & Fishback, P. V. (1999). The impact of institutional change on compensating wage differentials for accident risk: South korea, 1984–1990. *Journal of Risk and Uncertainty*, *18*, 231–248.

Kremer, M., & Levy, D. (2008). Peer effects and alcohol use among college students. *Journal of Economic Perspectives*, *22*(3), 189–206.

Kuchler, T., & Zafar, B. (2019). Personal experiences and expectations about aggregate outcomes. *The Journal of Finance*, *74*(5), 2491–2542.

Lavetti, K. (2020). The estimation of compensating wage differentials: Lessons from the deadliest catch. *Journal of Business & Economic Statistics*, *38*(1), 165–182.

Lavetti, K., & Schmutte, I. M. (2023). Gender differences in sorting on wages and risk. *Journal of Econometrics*, *233*(2), 507–523.

Le Barbanchon, T., Rathelot, R., & Roulet, A. (2021). Gender differences in job search: Trading off commute against wage. *The Quarterly Journal of Economics*, *136*(1), 381–426.

Leeth, J. D., & Ruser, J. (2003). Compensating wage differentials for fatal and nonfatal injury risk by gender and race. *Journal of Risk and Uncertainty*, *27*, 257–277.

Lieber, E. M., & Skimmyhorn, W. (2018). Peer effects in financial decision-making. *Journal of Public Economics*, *163*, 37–59.

Ling, H., Hardy, J., & Zetterberg, H. (2015). Neurological consequences of traumatic brain injuries in sports. *Molecular and Cellular Neuroscience*, *66*, 114–122.

Malmendier, U., & Nagel, S. (2011). Depression babies: Do macroeconomic experiences affect risk taking? *The Quarterly Journal of Economics*, *126*(1), 373–416.

Manski, C. F. (1993). Identification of endogenous social effects: The reflection problem. *The Review of Economic Studies*, *60*(3), 531–542.

Marmaros, D., & Sacerdote, B. (2002). Peer and social networks in job search. *European Economic Review*, *46*(4-5), 870–879.

Martland, H. S. (1928). Punch drunk. *Journal of the American Medical Association*, *91*(15), 1103–1107.

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

McCoy, S. J., & Walsh, R. P. (2018). Wildfire risk, salience & housing demand. *Journal of Environmental Economics and Management*, *91*, 203–228.

McKee, A. C., Stein, T. D., Kiernan, P. T., & Alvarez, V. E. (2015). The neuropathology of chronic traumatic encephalopathy. *Brain Pathology*, *25*(3), 350–364.

Michaelides, M. (2010). A new test of compensating differences: Evidence on the importance of unobserved heterogeneity. *Journal of Sports Economics*, *11*(5), 475–495.

Mridha, H., & Khan, F. (2013). Does unemployment affect compensating wage differentials for work-related fatal injuries? Empirical evidence from occupational data. *Applied Economics Letters*, *20*(14), 1283–1287.

Nagler, M., Rincke, J., & Winkler, E. (2023). High-pressure, high-paying jobs? *Review of Economics* *and Statistics*, 1–45.

Omalu, B. I., DeKosky, S. T., Hamilton, R. L., Minster, R. L., Kamboh, M. I., Shakir, A. M., & Wecht, C. H. (2006). Chronic traumatic encephalopathy in a National Football League player: Part II. *Neurosurgery*, *59*(5), 1086–1093.

Omalu, B. I., DeKosky, S. T., Minster, R. L., Kamboh, M. I., Hamilton, R. L., & Wecht, C. H. (2005). Chronic traumatic encephalopathy in a national football league player. *Neurosurgery*, *57*(1), 128–134. https://doi.org/10.1227/01.neu.0000163407.92769.ed

Oster, E., Shoulson, I., & Dorsey, E. R. (2013). Optimal expectations and limited medical testing: Evidence from huntington disease. *American Economic Review*, *103*(2), 804–830.

Oster, E., & Thornton, R. (2012). Determinants of technology adoption: Peer effects in menstrual cup take-up. *Journal of the European Economic Association*, *10*(6), 1263–1293.

Quinn, J. F. (1977). Microeconomic determinants of early retirement: A cross-sectional view of white married men. *Journal of Human Resources*, *12*(3), 329–346.

Robertson, L. S., O'Neill, B., & Wixom, C. W. (1972). Factors associated with observed safety belt use. *Journal of Health and Social Behavior*, 18–24.

Rosaz, J., Slonim, R., & Villeval, M. C. (2016). Quitting and peer effects at work. *Labour Economics*, *39*, 55–67.

Rosen, S. (1981). The economics of superstars. *The American Economic Review*, *71*(5), 845–858.

Rosen, S. (1986). The theory of equalizing differences. *Handbook of Labor Economics*, *1*, 641–692.

Sacerdote, B. (2001). Peer effects with random assignment: Results for Dartmouth roommates. *The Quarterly Journal of Economics*, *116*(2), 681–704.

Schildberg-Hörisch, H. (2018). Are risk preferences stable? *Journal of Economic Perspectives*, *32*(2), 135–154.

Shilts, R. (1987). *And the band played on*. St. Martin's Press.

Stern, R. (2015). *Study of retirees links youth football to brain problems*. The New York Times. https://www.nytimes.com/2015/01/29/sports/football/study-points-to-cognitive-dangers-of-tackle-football-before-age-12.html

Stevenson, M. (2017). Breaking bad: Mechanisms of social influence and the path to criminality in juvenile jails. *Review of Economics and Statistics*, *99*(5), 824–838.

Viscusi, W. K. (1980). Imperfect job risk information and optimal workmen's compensation benefits. *Journal of Public Economics*, *14*(3), 319–337.

Weinstein, N. D. (1989). Optimistic biases about personal risks. *Science*, *246*(4935), 1232–1233.

Yakusheva, O., Kapinos, K. A., & Eisenberg, D. (2014). Estimating heterogeneous and hierarchical peer effects on body weight using roommate assignments as a natural experiment. *Journal of Human Resources*, *49*(1), 234–261.

# Appendix

## Appendix Table of Contents

## A1 Worker Productivity

Figure A1 displays a collection of descriptive figures related to worker productivity in professional football and hockey. Figure A1a plots the relationship between age and productivity, following an inverse 'U' shape. Notably, the upward slope during an athlete's early years is steeper than the subsequent decline, consistent with the idea that the rate of human capital acquisition initially outpaces the effects of physical decline, which typically begins in the early to mid-thirties. Figure A1b shows the relationship between productivity and award recognition. While the average productivity score is around 2, this measure increases significantly, reaching 9–11 for athletes named to All-Star or Pro Bowl teams. For football players selected to the All-Pro team, average productivity exceeds 12. The productivity metric in this figure is presented in its unadjusted form.

Figure A1c shows the relationship between average career productivity and career length. Productivity is a strong predictor of career duration, as indicated by the steep curves, which suggest an estimated four additional years of career length for every half-standard deviation increase in average productivity. Notably, players with careers spanning ten or more seasons demonstrate disproportionately higher productivity levels. Figure A1d displays hows the relationship between productivity and the number of years relative to a player's eventual exit from professional sports. Here, productivity declines sharply as players approach the end of their careers, indicating that most exits from sports are involuntary and underscoring the significant role of injury in retirement decisions.

Finally, figure A1e llustrates the distributions of worker productivity and salary across hockey and football. These measures are centered at zero by design, approximating the value of a player in a given year relative to what would be expected from a replacement or "bench" player. Both distributions are right-skewed, indicating a small percentage of players who have outsized impacts on team success. However, football exhibits a much longer right tail in the productivity distribution, whereas hockey has a larger concentration of players with productivity values at or near zero. This suggests that superstars are more prevalent in football, likely due to rules that favor offensive playmaking and structural factors like available space. For example, though football has nearly twice as many players on the field (11) as hockey has on the ice (6), football players enjoy nearly double the available space per player, with fields measuring approximately 57,000 square feet compared to hockey rinks at around 17,000 square feet. These spatial and structural differences may be a key factor in football's tendency to generate standout individual performers.

Table A1 presents the highest career productivity totals in the NFL, measured by 'Approximate Value.' Leading the list is Tom Brady, widely regarded as the greatest football player of all time, having won an unprecedented seven Super Bowl championships during an era of (rela-

**Figure A1.** Trends and Distributions in Worker Productivity

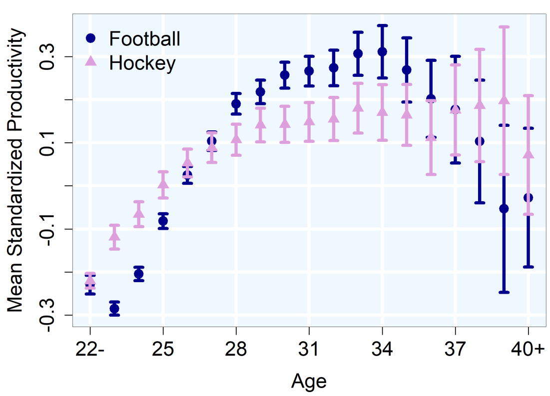


**(a)** Worker Productivity & Age

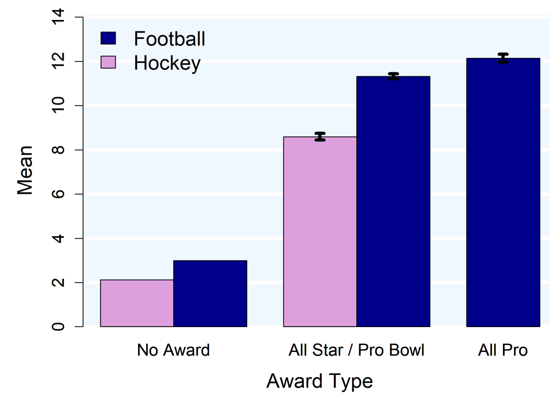


**(b)** Productivity by Award Type

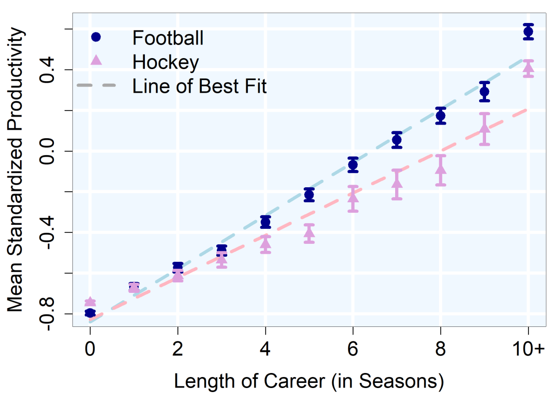


**(c)** Productivity & Career Length

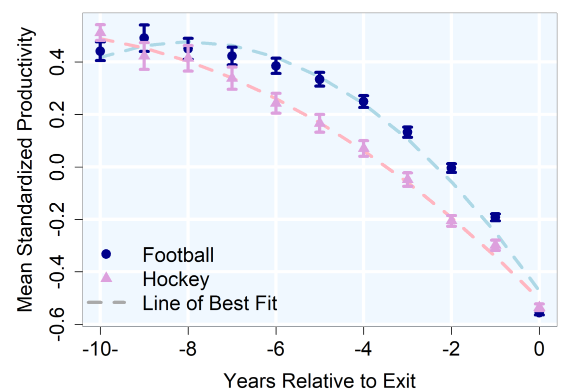


**(d)** Productivity Relative to Exit

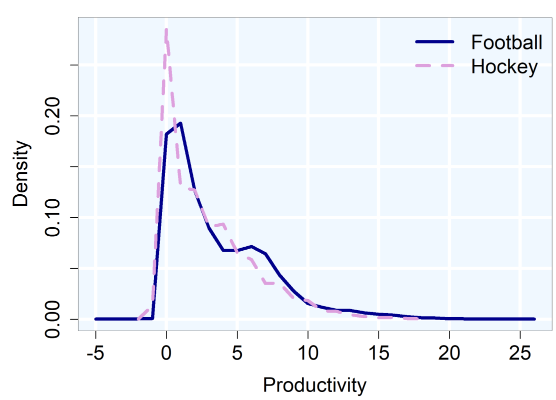


**(e)** Productivity Distribution

tive) competitive parity in the NFL. Brady not only holds the highest total approximate value in NFL history but also ranks among the top five in average value per season. Legendary defensive lineman Aaron Donald, with the highest all-time approximate value per season, appears only in the top 75 overall due to his early retirement; he nonetheless achieved Pro Bowl selections every season of his career and was named to the All-Pro team in all but two seasons.

**Table A1.** Highest Career Productivity Totals in the NFL

| Ranking | Player | Position | Total AV | Years | Seasons | AV/Season |
|---|---|---|---|---|---|---|
| 1 | Tom Brady | QB | 326 | 2000-2022 | 22 | 14.8 |
| 2 | Drew Brees | QB | 277 | 2001-2020 | 19 | 14.6 |
| 3 | Peyton Manning | QB | 271 | 1998-2015 | 17 | 15.9 |
| 4 | Brett Favre | QB | 259 | 1991-2010 | 19 | 13.6 |
| 5 | Jerry Rice | WR | 251 | 1985-2004 | 19 | 13.2 |
| 6 | Fran Tarkenton | QB | 233 | 1961-1978 | 17 | 13.7 |
| 7 | Aaron Rodgers | QB | 231 | 2005-2023 | 18 | 12.8 |
| 7 | Reggie White | DE | 231 | 1985-2000 | 15 | 15.4 |
| 9 | Bruce Smith | DE | 229 | 1985-2003 | 18 | 12.7 |
| 10 | Ray Lewis | LB | 224 | 1996-2012 | 16 | 14.0 |
| ... | | | | | | |
| 18 | Junior Seau | LB | 195 | 1990-2009 | 19 | 10.2 |
| ... | | | | | | |
| 60 | Aaron Donald | DT | 153 | 2014-2023 | 9 | 17 |

This table presents a ranked list of professional football players who contributed most to their teams' success, as measured by "Approximate Value" (AV), a metric used here to capture player productivity. "Total AV" reflects the cumulative AV over a player's career, while "AV/Season" represents the average AV per season. Position abbreviations include "QB" for quarterback, "WR" for wide receiver, "DE" for defensive end, "LB" for linebacker, and "DT" for defensive tackle. The "Years" column shows the years they were active, and "Seasons" indicates the total number of seasons played in professional football.

Notably, six of the top seven players by this measure are quarterbacks, underscoring the validity of approximate value as a productivity measure and highlighting the widely held view of quarterback as the 'most difficult and important position in sports.' This also reinforces the pattern shown in Figure A1e that the structure of football favors offensive production and contributes to the creation of superstars.

Hall of Fame linebacker Junior Seau, ranking in the top 20, remains one of the most poignant figures in the history of CTE awareness. His death by suicide in 2012 sent shockwaves through the sports world and marked a turning point in discussions about CTE. Seau reportedly took deliberate steps to preserve his brain for research, allowing for the posthumous diagnosis of CTE, which was confirmed shortly thereafter.

Table A2 presents the highest career productivity totals in the NHL, as measured by "Point Shares". At the top of the list is Wayne Gretzky, widely recognized as the greatest hockey player of all time. Not only does Gretzky hold the highest total point shares in NHL history, but he also has the second-highest average point shares per season, surpassed only by Connor McDavid. McDavid, often considered the greatest player of his generation, is the most recent recipient of the NHL's Most Valuable Player award and is on track to enter the top 100 all-time point shares before completing his tenth professional season. Legendary defenseman Bobby Orr ranks in the top five for average point shares per season but only appears in the top 50 for total point shares due to a career shortened by knee injuries. All players listed in this table who are not still actively playing have been inducted into the NHL Hall of Fame.

**Table A2.** Highest Career Productivity Totals in the NHL

| Ranking | Player | Position | Total PS | Years | Seasons | Mean PS |
|---|---|---|---|---|---|---|
| 1 | Wayne Gretzky | C | 251.01 | 1979-99 | 20 | 12.6 |
| 2 | Ray Bourque | D | 242.69 | 1979-01 | 22 | 11.0 |
| 3 | Roberto Luongo | G | 217.84 | 1999-19 | 20 | 10.9 |
| 4 | Gordie Howe | RW | 217.11 | 1946-80 | 34 | 6.4 |
| 5 | Jaromír Jágr | RW | 217.06 | 1990-18 | 28 | 7.8 |
| 6 | Nicklas Lidström | D | 211.77 | 1991-12 | 21 | 10.1 |
| 7 | Martin Brodeur | G | 206.97 | 1991-15 | 24 | 8.6 |
| 8 | Alex Ovechkin | LW | 203.68 | 2005- | 19 | 10.7 |
| 9 | Patrick Roy | G | 198.34 | 1984-03 | 19 | 10.4 |
| 10 | Al MacInnis | D | 195.01 | 1981-04 | 23 | 8.5 |
| ... | | | | | | |
| 42 | Bobby Orr | D | 150.95 | 1966-79 | 13 | 11.6 |
| ... | | | | | | |
| 129 | Connor McDavid | C | 114.76 | 2015- | 9 | 12.8 |

This table provides a ranked list of professional hockey players with the highest career productivity, as measured by "Point Shares" (PS), which captures a player's contribution to team success. "Total PS" represents the cumulative point shares over each player's career, while "Mean PS" indicates the average point shares per season. Positions are abbreviated as follows: "C" for center, "RW" for right wing, "LW" for left wing, "D" for defenseman, and "G" for goaltender. The "Years" column shows the active years of each player, while "Seasons" reflects the total seasons played in the NHL.

**Click to Return to Appendix Table of Contents**

## A2 Injuries

Football teams are required to disclose the injury status of players. I gather data on the full range of disclosed injuries from prosportstransactions.com and analyze the text of these injury announcements to categorize each player's injury type. Figure A2 presents a 3x1 layout of these data as categorical time series. Figure A2a shows the proportion of games missed due to injury in each season. Lower body injuries are by far the most common, accounting for approximately 60% of all missed games. Injury measures have remained relatively stable over time, with the notable exception of head and neck injuries, which rose from 10% of all missed games in 2000 to 20% in 2020. This increase is particularly striking, as players with concussions were seldom withheld from games for more than one week, while common lower body injuries, such as ACL sprains or tears, typically require rehabilitation periods ranging from 2–10 months.

**Figure A2.** Injuries in the NFL over Time

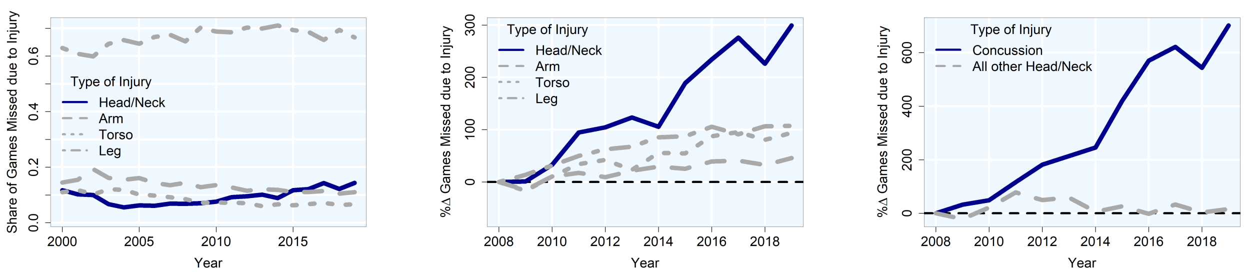


**(a)** Proportion of Injuries Types **(b)** Change in Injuries by Type **(c)** Change in Head Injuries

Figure A2b shows that from 2008 to 2020, all injury types have shown increases, with three of the four types doubling in frequency, reflecting shifts in protocol and attitudes toward health and safety. Head injuries, in particular, have increased at a much higher rate due to changes in NFL concussion protocols introduced in the 2009 and 2010 seasons. These protocols mandated that any player diagnosed with a concussion be removed from play and prohibited from returning the same day if they showed symptoms like memory loss or loss of consciousness. Additionally, the league encouraged teams to consult independent neurological specialists, required players to be fully asymptomatic and cleared by medical professionals before returning, and began imposing stricter penalties for illegal head-to-head hits. Figure A2 further breaks down head and neck injuries into concussions and non-concussions, showing a 600% increase in games missed due to concussions. This accounts for the entire rise in head and neck injuries documented in Figure A2b.

**Click to Return to Appendix Table of Contents**

## A3 Worker Turnover

Figure A3a displays very high worker (athlete) turnover rates in professional sports – roughly 20% in football and 14% in hockey. Figure A3b shows that these turnover rates result in notable differences in career length; the median career spans four years in hockey compared to three years in football. With the exception of a 2005 labor dispute among hockey players, exit rates in these professions have remained remarkably stable over time.

**Figure A3.** Trends and Distributions of Career Longevity across Sports

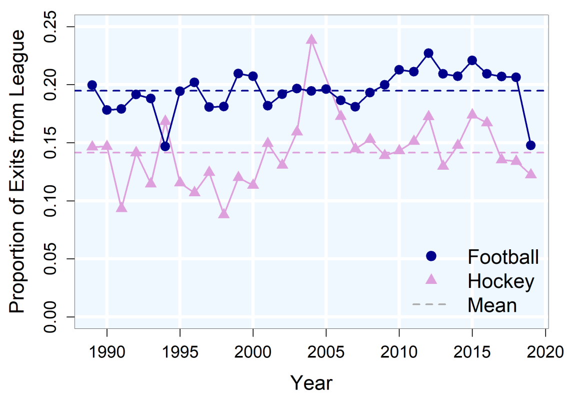


**(a)** Timeseries of Worker Turnover

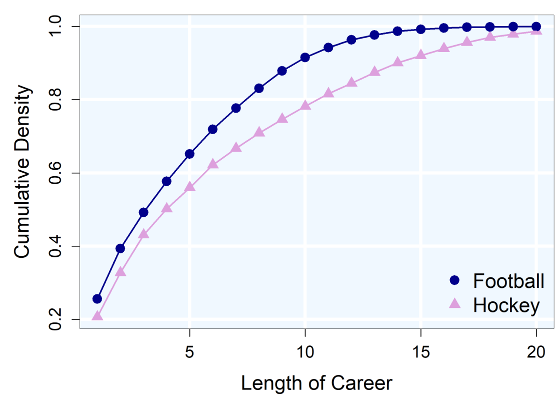


**(b)** Distribution of Career Longevity

Large differences in the turnover rate has implications for the the identification of peer effects in this setting. Workers "treated" by the death of a former colleague are likely to be significantly older and more productive than typical non- or not-yet-treated workers since "treatment" is likely to have a time lag between when individuals were last teammates and their death. Visual confirmation of concerns that using the full-sample of non- and not-yet-treated workers as a control group is provided in Figure A4 which shows substantial differences in exit rates between treated and untreated workers leading up to the CTE death of a former teammate.

**Figure A4.** Unadjusted Cumulative Share of Cohort Exits

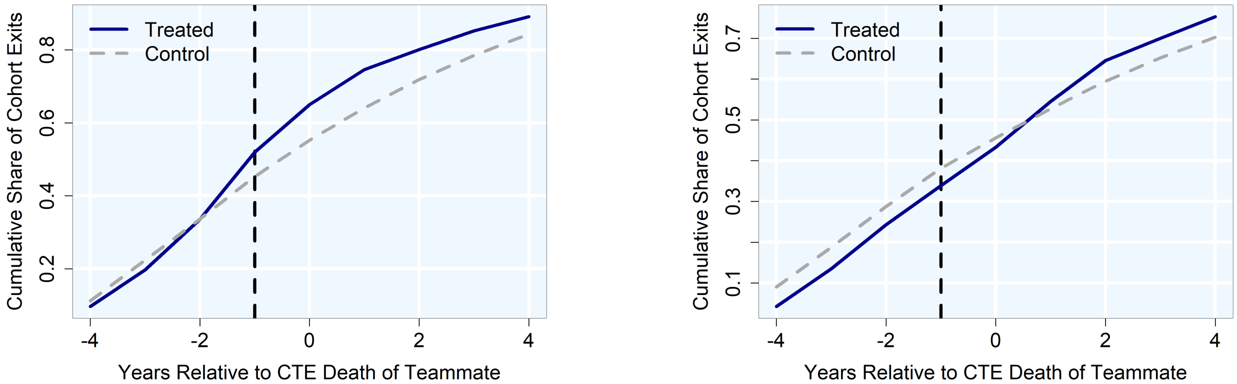


**(a)** Football **(b)** Hockey

**Click to Return to Appendix Table of Contents**

## A4 Viewership of Hockey and Football

I gather and analyze two sources of data related to sports popularity and how this influences popular awareness of CTE deaths within the sample period of analysis for this paper, with football far outpacing hockey in both cases.

The first source of data is gathered from Nielsen Media Research estimates of viewership for the championship game(s) for the NHL and NFL. Figure A5 reveals that viewership of American football far outpaces that of hockey, regularly gaining 20-40 times the viewership. The viewership of the gridiron football has declined slowly in the past decade, but not substantially. Some have attributed this decline to the growing awareness of CTE, which diminishes the enjoyment of watching the games. The discrepancy in attention paid to these sports likely does not only apply to the championship games, as recent survey evidence has confirmed that NFL regular season games now regularly make up 90+% of television programs watched each year in the United States (Adgate, 2024).

**Figure A5.** Neilsen Ratings for Champsionship Game / Series

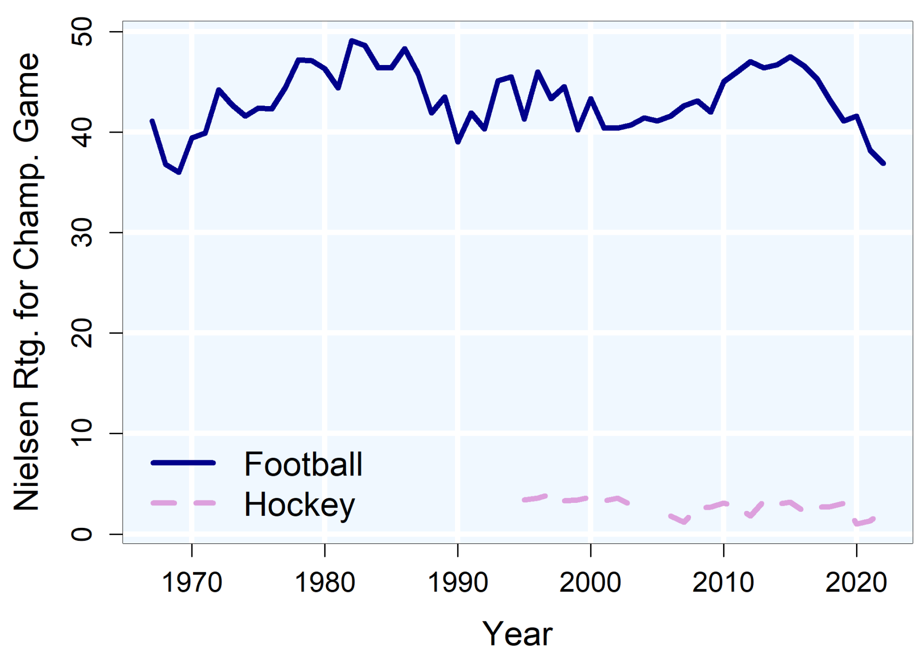


**Click to Return to Appendix Table of Contents**

## A5 Public Awareness of Athlete CTE Diagnoses

I then gather data from Google Trends, which enables me to track changes in internet search intensity for topics related to CTE. This data, shown in Figure A6, aligns with the history of CTE entering the public vernacular following the deaths of four hockey players in 2010 and 2011. The first significant increase in search interest appears in 2012, corresponding with the suicide of legendary football player Junior Seau, who reportedly left a note requesting that his brain be examined for CTE.

**Figure A6.** Timeline of Variation in Google Search Intensity for "CTE"

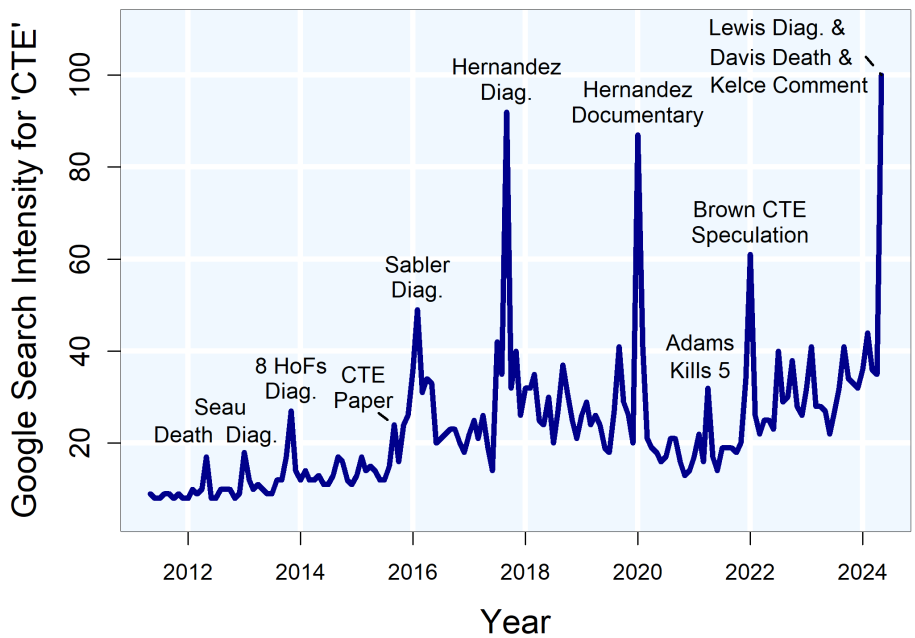


The second noticeable increase in search intensity corresponded to reports that eight Hall of Fame football players had either been formally diagnosed with CTE or were showing signs of significant cognitive decline. This increase included coverage of the health struggles of beloved former Dallas Cowboys running back Tony Dorsett. A subsequent rise in search intensity followed the release of a case study examining the brain tissue of nearly 80 former NFL players, accompanied by op-eds in national outlets like The New York Times that linked CTE to youth football participation (McKee et al., 2015; Stern, 2015). However, this increase was eclipsed by the surge in search interest upon the news that former quarterback Ken Stabler, a recipient of the prestigious MVP award during his playing years, had been diagnosed with the disease.

The most significant increases in search interest for CTE centered around Aaron Hernandez, a high-profile collegiate and NFL player whose case drew national attention after his murder conviction. Following his death by suicide in prison, Hernandez was posthumously diagnosed with CTE. A 2020 Netflix documentary chronicling his life underscored the potential influence of CTE on his behavior, sparking further public awareness about the disease. Another notable case was the violent death of former NFL player Phillip Adams, who killed five people in a 2021 shooting. Interestingly, the spike in CTE-related searches occurred upon news of his death rather than his diagnosis, suggesting that public understanding of the link between aggression and CTE had deepened following the release of Hernandez's documentary.

More recent spikes in search intensity followed the in-game retirement of star wide receiver Antonio Brown. Brown's erratic actions and increasingly frequent emotional outbursts became a topic of online speculation, with many attributing his behavior to possible CTE symptoms. The largest and most recent surge in search interest occurred in the summer of 2024, amid three back-to-back events: the CTE diagnosis of NFL Hall of Famer Ray Lewis' 28-year-old son, the sudden death of former player Vontae Davis at age 35—known for retiring mid-game—and a public admission from retired star lineman Jason Kelce, who expressed concerns that he might have CTE.

Figure A6 highlights a notable pattern: each spike in search intensity aligns with the death of a football player, rather than a hockey player. To assess whether hockey player deaths serve as significant information shocks to the general public, I analyze changes in internet search intensity before and after CTE diagnoses. Specifically, I convert the search intensity metric (ranging from zero to one hundred) into a percentage change format to ensure that earlier periods, when CTE was less widely recognized, are weighted equally to more recent periods. I then "stack" each athlete's CTE diagnosis into an event-time window of plus or minus seven months. By including an indicator for the timing of each athlete's death, I estimate changes in search intensity relative to the months leading up to each event. Note that due to the high frequency of CTE diagnoses, there is no "untreated" control group of months.

The analysis results are presented in Table A3. Model 1 shows an approximate 15 percentage point increase in search intensity for CTE following the death of an athlete later diagnosed with the disease. This effect size reduces slightly to 13 percentage points, remaining statistically significant when controlling for month and year fixed effects, indicating that seasonal factors, such as league play cycles, do not explain the increase. Models 3 and 4 further differentiate the impacts by the athlete's sport. As expected, the effect is larger for football, with a 16 percentage point increase. However, hockey also shows a substantial effect, with a 10 percentage point increase. These findings suggest that while football-related deaths have a more pronounced impact on CTE awareness, hockey-related deaths also contribute significantly, though to a lesser extent.

Figure A7 displays that the results from A3 in an event study format by aggregating the outcome variable – internet search intensity for "CTE – following the death of an athlete later diag-

**Table A3.** The Effect CTE Deaths on Internet Search Intensity

| Dependent Variable: | Search | | | |
|---|---|---|---|---|
| Sample: | Full | | Hockey | Football |
| Model: | (1) | (2) | (3) | (4) |
| CTE (Post) | 0.15*** | 0.13** | 0.10** | 0.16*** |
| | (0.05) | (0.05) | (0.05) | (0.06) |
| *Fixed-effects* | | | | |
| Month | | ✓ | | |
| Year | | ✓ | | |
| Observations | 248 | 248 | 88 | 205 |

*IID standard-errors in parentheses.*

nosed with the disease. The outcome is transformed by dividing the value at period $t - 1$ and subtracting one, allowing it to be interpreted as changes in search intensity. The figure reveals a slight decrease in search intensity, approximately 10%, in the period leading up to these deaths, followed by an asymmetrically larger increase of 13-18% afterward, suggesting that these deaths helped sustain media attention on the disease.

**Figure A7.** Event Study of CTE Deaths on Internet Search Intensity

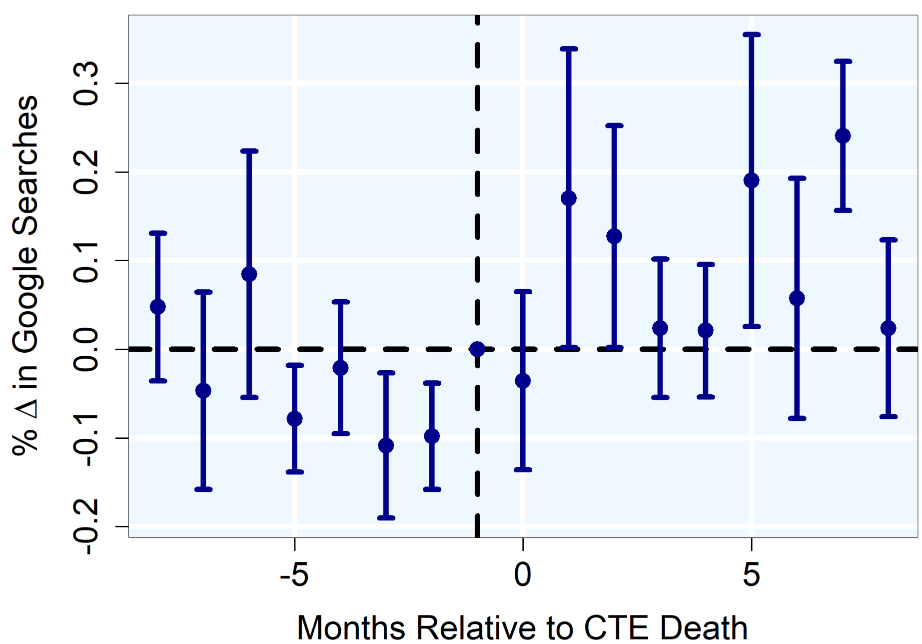


[Click to Return to Appendix Table of Contents]

## A6 Teammate Characteristics

Figure A8b displays the differences in characteristics of "exposure" to colleagues who went on to be diagnosed with CTE. Figure A8a displays the amount of time that has elapsed between the period when individuals last were on the same team and the date of the death of their former colleague. 5 years have passed since last teammates for the median "treated" athlete in my sample, though there remains large variation, ranging from 0 to 19 years. Figure A8b displays the amount of time that athletes spent on the same team at any point in their career. The median athletes were teammates for 2 seasons. Though there is a tremendous amount of variation in this measure, ranging from 1 to 14 years, fewer than 10% of athletes have been teammates with someone who went on to die with CTE for more than 4 years.

**Figure A8.** Distribution of Time as (and since) Teammates

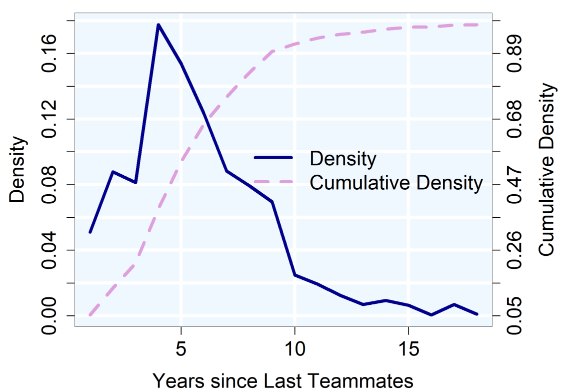


**(a)** Distribution of Years since Last Teammates

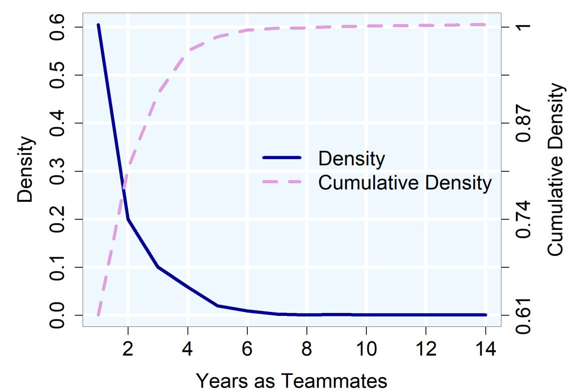


**(b)** Distribution of Years as Teammates

**Click to Return to Appendix Table of Contents**

## A7 Salaries & Contracts

A common misconception about differences between football and hockey players are differences in their annual earnings. While large differences are clearly visible in average compensation, this masks large similarities in the distribution of these earnings. For instance, as shown in Figure A9 the median hockey player has marginally greater earnings than the same for football. Differences in earnings emerge at 99th percentile of the earnings distribution where football players earn more than hockey players by a factor of 4.

**Figure A9.** Annual Salary Distribution

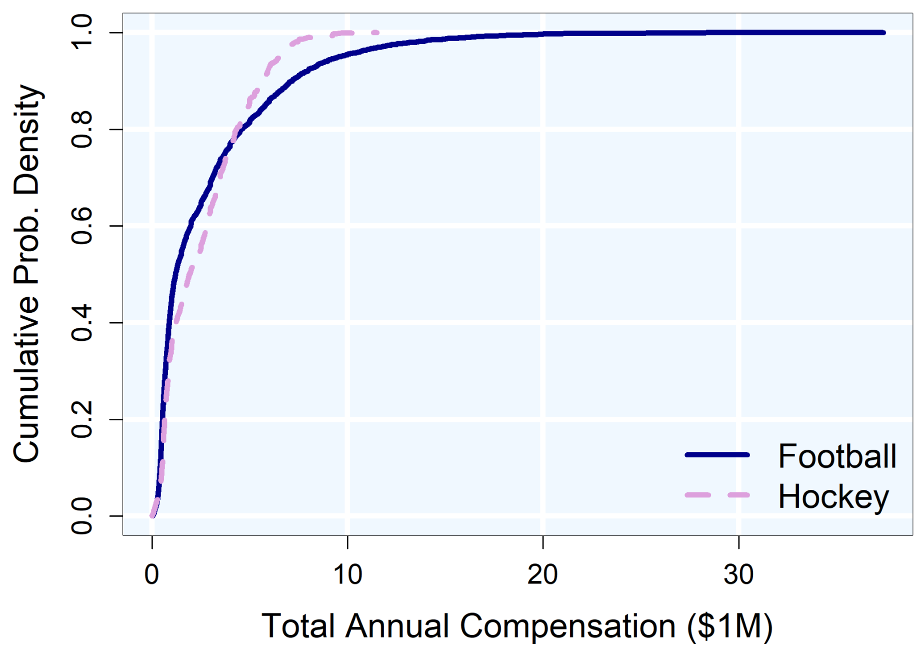


However, Figure A9 should be interpreted with a degree of caution as I do not have access to the full universe of players' contract data as shown in Figure A10. Figure A10b displays the trends in the share of the population of all athletes for whom I cannot observe the details of their contracts across sports. Beginning in 2005, I observe contract data for approximately 25% of football players and fewer than 1 percent of hockey players. These rise to approximately 85% by 2020. However, I observe roughly 40% of all contracts at the start of my analysis period in both sports as CTE in hockey players was discovered later than for football.

Figure A10a reveals, unsurprisingly, that there is non-randomness in whose contracts are observable. Athletes with full contract information are over 20% as as productive in both sports. For football players with contract data, they have career lengths which are 30% longer though this difference is insignificantly different for hockey.

**Figure A10.** Trends in Contract Missingess

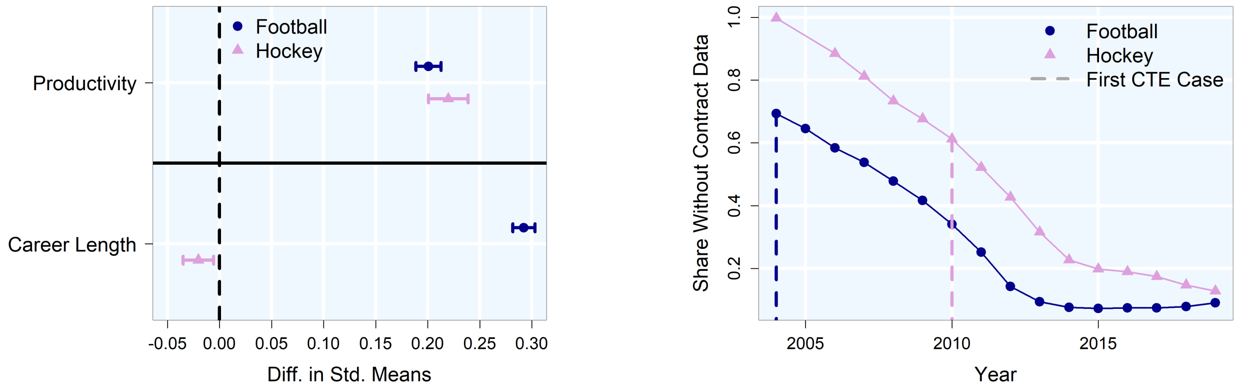


**(a)** Predictors of Contract Data

**(b)** Share without Contract Data

**Click to Return to Appendix Table of Contents**

## A8 Treatment Stratification

Throughout the paper, there are numerous instances where the treated group is stratified by degrees of "exposure" to former teammates who were diagnosed with CTE postmortem. One empirical concern with changing the construction of the treated group is that it reduces its similarity to the control group, harming identification and the research design. The most obvious analysis in the main body of the paper where this is likely to have occurred is in panel B of Table 5 which stratifies the main result by age, a highly influential variable with respect to retirement decisions of athletes.

As stated within the body of the paper, I preform a coarsened exact matching (CEM) method using the variables of age and productivity for altering the control group each time the treated group is stratified. In order to provide visual evidence of the efficacy of this method, Figure A11 presents the event study estimates from panel B of Table 5. Though each age bin in this event study has vastly different rates of exit from the profession, the matching method successfully omits a sufficient number of "bad" observations from the control group, achieving parallel pre-trends.

**Figure A11.** Event Study of the Effect of the CTE Death of a Former Teammate on Professional Exit; Stratified by Age at Time of Treatment

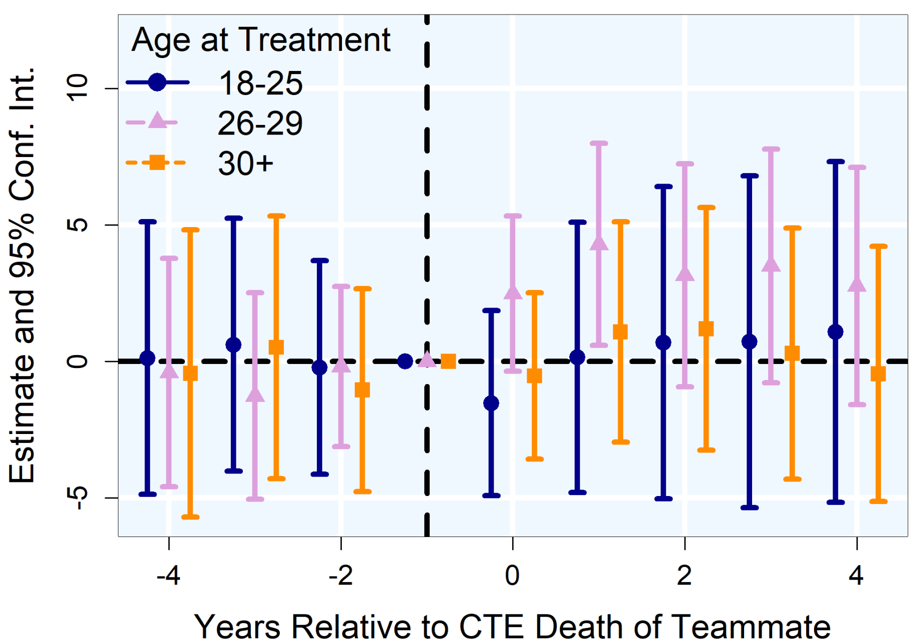


**Click to Return to Appendix Table of Contents**