

Chemistry Under Pressure: Measuring NBA Team Resilience Through Network Structure

Aakanksh Vaidya, Manraaj Singh, Daniel Wurzburger

Course: QSS 41

Date: 8th June 2025

Abstract

This study investigates how chemistry distributions affect team resilience to injury in the NBA. We crafted a metric, Impact Per Minute (IPM), capturing positive and negative contributions in player pairings for each NBA team, and quantified team “spikiness” with the coefficient of variance. We hypothesized that teams with uniformly distributed IPM networks would face a smaller IPM-penalty after shocks (injuries) than spiky networks. However, we found no significant difference ($p = 0.73$). Instead, we found an extreme difference in the volatility of injury’s impact ($\sigma = 6.92$ vs $\sigma = 21.35$, $p < 0.0001$). These findings have practical applications for front office decisions in the NBA. A playoff quality team may benefit most from evening out its roster, acquiring or trading for players that will increase redundancy between IPM connections. On the other end of the spectrum, a rebuilding team with a longshot at making the playoffs may benefit from acquiring a star, as their low upside makes it necessary to embrace risk for any chance of success.

1. Introduction

Team Chemistry is a concept often discussed in sports but it's rarely quantified with empirical rigor. It's an elusive quality, intuitively understood as the synergy between players that goes beyond their individual performances. This analysis will move beyond anecdotal evidence to quantitatively measure team chemistry and its effect on performance, in particular when a team faces adversity like a significant player injury.

In the context of the NBA, where the contributions of star players are immense and player rotations are constant, a team's structure is always put to a test. While built-in breaks throughout a 48-minute game for key players function as predictable, micro-level disruptions to a team's momentary flow (though often beneficial), a major, prolonged injury represents a much larger, unforeseen shock to the team system.

This project explores how the underlying structure of a team's chemistry network influences its ability to withstand such shocks. We ask: Do teams that rely on a few elite player pairings suffer more from an injury than teams where performance is more evenly distributed across the roster? To answer this, we developed a network-based approach to model team chemistry and analyze its performance under stress.

2. Literature review

2.1 Mukherjee et al. (2018)

One attempt at quantifying the impact of chemistry can be seen in Mukherjee et al. (2018). They used a naive measure of chemistry – simply, the amount of past wins players have had together. They did this for the NBA, MLB, EPL, IPL, and even Dota 2 (a video game). In baseball, for

example, they controlled for estimated talent (using an average of hitter and pitcher stats from the prior 10 years) as well as team fixed effects, and found that teams with higher team chemistry tended to score more than expected to a statistically significant degree ($P < .001$). Including chemistry instead of just team skill increased the model's predictive power: the R^2 for predicting run differential in individual games increased from 6.4 to 10.5%. Fascinatingly, after accounting for team chemistry, the team skill estimators (past WAR / OPS) did not influence the score in a statistically significant manner. This brings up potential errors with the interpretation of the study, particularly its chemistry metric. It is quite possible that the chemistry metric is capturing something different than is intended – maybe, players with shared high win totals who are still on the same team have done so due to long-term expensive contracts, implying that they project better than a simple averaging of past WAR / OPS might indicate. The study could've accounted for this by using player projections to more accurately estimate team strength. If chemistry was shown to have a predictive effect using forward-looking projections instead of backwards-looking performance, there would be a more convincing argument for it as a legitimate effect. Still, the fact that the study found a statistically significant chemistry effect across such different forms of competition – MLB, NBA, EPL, IPL and Dota 2 (p values all much less than 0.05) – provides compelling, if not definitive, evidence for its existence. It is unlikely that the methodological flaws would happen to produce statistically significant positive results across all of the diverse domains measured.

2.2 Pradeep (2024)

Pradeep (2024) focused exclusively on basketball and employed a more sophisticated and basketball-specific definition of chemistry. The study used roster continuity, the standard deviation of salaries and minutes played, and various in-game statistics such as passes to construct its chemistry metric. While the study found team chemistry to have a statistically significant effect on winning percentage (p-value of 0.0302), it did not find it to be of any real-world significance. The model that used chemistry to predict winning percentage had an R^2 of 0.019, meaning chemistry had very little impact on team success. Taken at face value, this study suggests that chemistry is irrelevant when considering roster construction. However, this study's use of chemistry as the sole predictor of winning percentage obfuscates the genuine impact it may have. Chemistry could correlate with good or bad overall skill, different playstyles, coaching strategies – or any number of other factors. These other factors could be influencing winning percentage in ways not considered by this regression, and as such, it only establishes a correlation. The true causal impact of chemistry on winning percentage could be lower or higher than presented here.

2.3 Fewell et al. (2012)

These two studies show two different ways of looking at chemistry, both in its definition and in its impact. They demonstrate a larger theme of its murkiness – intuitively it can be understood as the ability of players to work well together, to synergize, but it is harder to pin down as an exact concept. A more grounded way to look at chemistry could be through structural stability. This

means chemistry is less focused on how well players work together on the court or on the field, but instead looks at the balance of the roster, and questions how durable it is as a unit. This may seem counterintuitive at first. Chemistry is often imagined as the combination of irreplaceable players who just click together – but this makes a roster less versatile. Fewell et al. (2012) helps frame this idea. The study modeled NBA offenses as passing graphs and found that more evenly distributed flow centrality correlated to higher playoff success in the 2010 postseason. While not fully causal, this supports the idea that evenly distributed ball movement structures are beneficial.

3. Hypothesis and Constructed Variable

A chemistry network that is evenly distributed should, in principle, possess greater resilience to the adverse effects of random shocks like player injuries. If each player pairing offers a similar level of on-court performance, the loss of one pairing can be more easily absorbed by others that provide comparable value. Therefore, our main hypothesis is that teams with more uniformly distributed two-player IPM networks will experience less performance volatility following a significant injury than teams with "spiky" (uneven) IPM distributions.

To test this, we attempted to capture on-court chemistry using our Impact Per Minute (IPM) metric, a composite statistic designed to capture the net effect of a player duo. The formula is:

$$\text{IPM}_{ij} = \frac{\text{PTS} + \text{FGM} + \text{FTM} - \text{FGA} - \text{FTA} + \text{DREB} + \frac{1}{2} \cdot \text{OREB} + \text{AST} + \text{STL} + \frac{1}{2} \cdot \text{BLK} - \text{PF} - \text{TO}}{\text{MIN}}$$

This formula aggregates positive contributions (scoring, playmaking, defense), penalizes inefficiency and mistakes, and normalizes the result by minutes played to allow for fair comparison across lineups. While stats like points can be generated individually, IPM emphasizes synergy through metrics like assists, which are inherently collaborative. However, we acknowledge its limitations; as a box-score-derived metric, IPM cannot capture critical off-ball actions like defensive rotations, setting screens, or a player's leadership, which are all real components of true chemistry. Despite this, IPM provides a fairly robust, pairwise, and minute-weighted assessment of a pairing's tangible effectiveness on the court.

4. Data Collection and Network Construction

4.1 Data Sources

The data for this study was retrieved from publicly available NBA statistics for the 2024-2025 regular season. A web scraper was developed using the Python Selenium library to systematically gather two-man lineup data for all 30 NBA teams. For each team, this process produced a distinct CSV file containing the average performance statistics and minutes played for every player pairing that shared time on the court.

4.2 Network Creation

To investigate the structure of each team's chemistry, we first translated their on-court data into a formal network model. For each team, we constructed a weighted, undirected network where individual players are the nodes. An edge was created between any two players who shared minutes on the court, and the edge weight was set to their corresponding Impact Per Minute (IPM) score, to quantify the dyad's performance.


4.3 Cleaning and Backboning

The raw network data contains noise; not all pairings are strategically significant, as some players share limited minutes, which can lead to unreliable IPM scores and edges. To isolate the most meaningful connections, we performed a network backboning procedure. We used the disparity filter, a thresholding method ideal for weighted networks. Unlike a global filter that might eliminate all but the very strongest ties, the disparity filter evaluates the significance of each edge within the local context of its connected nodes. This approach preserves connections that are vital to a specific player, even if they aren't globally dominant. By setting an appropriate significance level, we pruned statistically insignificant pairings, resulting in a "backbone" of each team's core chemistry structure.


4.4 Network Structure Categorization

With the backboned networks, we categorized their overall structure using the Coefficient of Variation (CV). The CV, which is the ratio of the standard deviation of edge weights to the mean, serves as our normalized measure of "spikiness". A low CV indicates a "uniform" team with balanced chemistry, while a high CV points to a "spiky" team heavily reliant on a few elite pairings. We also calculated the Gini coefficient for each network, which confirmed the structural classifications derived from the CV, increasing our confidence in our methods

Uniform



Spiky



5. What the Results Can Inform

This analysis is useful beyond its novelty, and could be used to help teams with their risk profiles. Risk and volatility are important aspects of roster building that can be overlooked by traditional value metrics. We separate this into two categories, identification and application.

5.1 Identification

This network analysis provides a practical framework for NBA front offices to diagnose roster instability. By modeling team chemistry as a network, it allows teams to quantify their structural volatility, or "spikiness", and identify high-risk dependencies. This diagnostic approach can pinpoint specific vulnerabilities, such as an over-reliance on a few elite, high-IPM pairings (a "star-based hole") or weaknesses tied to a crucial "bridge" player who connects different team units.


5.2 Application

Once a team's structural profile is identified, front offices can tailor their strategy to their specific competitive situation. For a contending team aiming to maximize stability, the insights suggest they should mitigate risk by building a more "uniform" roster, targeting players who distribute chemistry evenly to ensure the team is resilient to injuries. Conversely, a different strategy is appropriate for rebuilding teams. For a team with low expectations, the high-risk, high-reward volatility of a "spiky," star-centric model is a logical path forward. The potential for a superstar to elevate the team's fortunes outweighs the risk of an inconsistent performance, as a uniformly average roster is unlikely to break through to the next level. This approach allows front offices to tailor their roster construction and risk profile to their specific competitive situation, moving beyond simply acquiring talent to building a more strategically sound team.

6. Modeling Adversity with Structural Holes


Injuries in the NBA, especially those that last multiple weeks, require a restructuring of the network as you are temporarily removing a node. As a result, the effect of that removal will enable you to characterize the role of the node and the role of its position. This structural hole analysis will lead to different outcomes depending on what node is removed. We posit the removal of a high impact player who largely contributes to wins on a team that has uneven chemistry distributions would be considered a star based hole.

Spiky Chemistry Network — Star Node Marked for Removal



We also identify the presence of a structural hole that occurs when removing a 6th or 7th teammate who has the most crossover playtime with both starting and bench players. We say the removal of this type of player is classified as a bridge based structural hole as their removal cuts the graph into two separate subgraphs.

Bridge-Based Structural Hole — Linking Starters and Bench



The star-based structural hole is more relevant when looking at injuries, as meaningful chemistry analysis would result in making claims against wins and losses. This correlates better with star-based holes as the player removed contributed to the team's game outcome.

7. The Gap Between Ideal and Actual Chemistry Measures

In a theoretically perfect network of team chemistry, betweenness centrality would be the most powerful tool for identifying structural holes. It quantifies how often a player lies on the shortest path between two others, making it ideal for detecting bridge-like roles – players who facilitate flow between otherwise disconnected subgroups (e.g., starters and bench units). If a team's chemistry score fully captured all meaningful aspects of performance – on-ball actions, off-ball movement, spatial pressure, communication, leadership, and tactical awareness – then players with high betweenness would accurately represent those holding the team together structurally and functionally.

However, in our dataset, chemistry is proxied by Impact per Minute (IPM), a metric built from box-score statistics accrued during shared floor time. While IPM captures co-scoring, rebounding, and playmaking events, it does not reflect intangible or unrecorded behaviors such as setting effective screens, pulling defenders away from the ball, or organizing team defense. Additionally, the co-scoring metrics aren't reflective of team-based plays as they are the sum of the metrics scored during shared court time and not necessarily due to teamwork. This creates a disconnect: edges in the graph are defined by visible statistical co-performance, not necessarily by the nuanced, often invisible dynamics that make certain players central to team cohesion.

As a result, when betweenness centrality is calculated using IPM as the edge weight, it tends to highlight players who log minutes across many lineup combinations, rather than those who strategically bridge structural gaps in chemistry. For example, backup guards or hybrid forwards may have high betweenness scores not because they orchestrate connectivity, but simply because they rotate in frequently. This weakens the signal that betweenness centrality was designed to capture.

We therefore conclude that while betweenness is conceptually ideal for identifying structural vulnerabilities, in practice it can only fulfill this role when the underlying chemistry measure is rich enough to reflect the true structural contributions of each player. Our current findings show how structural insights must be interpreted in light of measurement limitations – and how theoretical models sometimes outpace available data.


8. Data Analysis (Team-Level)

To ground our study in tangible examples, we conducted a case study analysis. First, we ranked all 30 NBA teams according to their "spikiness" using the Coefficient of Variation (CV) of their IPM networks. From this ranking, we selected two of the most "spiky" teams (the Los Angeles Clippers and Charlotte Hornets) and two of the most "uniform" teams (the Minnesota Timberwolves and Houston Rockets). For each of these teams, we identified a period of 3-4 weeks or more during the regular season where a player with a high average IPM sustained an injury.

We then collected two sets of lineup data: one from the stretch of games immediately before the injury, and another from the games during the player's absence. While generating new network graphs for each period was useful, the visual differences were difficult to interpret directly. To provide a clearer picture of the network's change, we calculated the percentage change in IPM for every player pairing from the "before" to the "during" period. We then visualized these results in the heatmaps below, where the color and intensity of each cell represent the magnitude and direction of the change in a duo's chemistry.

Spiky Teams


% change in IPM (after vs before)
NA = pair missing in one file or filtered out



1. Clippers

Following the injury, the Clippers' network showed a widespread, systemic decline in performance. Most remaining player duos experienced a significant drop in IPM, with changes frequently in the -12% to -22% range. This pattern suggests the "spiky" structure where chemistry was highly concentrated around the injured star. Once that central node was removed, the remaining connections weakened and the overall structure collapsed, as no other players were able to compensate for the loss. A few minor improvements likely stemmed from players being freed from that presence.

% change in IPM (after vs before)
NA = pair missing (or filtered) in one period




2. Hornets

The Hornets displayed a much more volatile and unpredictable response to losing a key player. The heatmap reveals a chaotic mix of outcomes: some pairings saw their performance drop sharply, while others experienced dramatic gains, with some improving by as much as 53%. This result is also implying a spiky structure, but one that is less fragile and more adaptable than the Clippers'. It suggests that while the team was still reliant on its star, his absence created opportunities for new, high-impact connections to form among the remaining players.


Uniform Teams:

% change in IPM (after vs before)
NA = pair missing (or filtered) in one period



1. Timberwolves

As the league's most uniform team, the Timberwolves demonstrated remarkable stability. The heatmap shows that most IPM changes were small and evenly dispersed across the roster, typically within a $\pm 10\%$ range. There were no terrible collapses or dramatic improvements; instead, the impact of the injury was absorbed by the network with minimal disruption. This muted response shows a well-balanced chemistry structure where the loss of a single player does not destabilize the entire system.



2. Rockets

The Rockets' provided an even stronger confirmation of our hypothesis. The heatmap is predominantly white, indicating that most IPM changes post-injury were clustered around 0%. This pattern of minimal disruption across the board shows a really deeply uniform chemistry structure. The mild and evenly distributed changes show a team whose performance is not contingent on any single player, allowing it to maintain its structural integrity and performance level even when facing adversity.

9. Difference in Means and Variance Test

To formally evaluate our hypothesis – that teams with spiky chemistry distributions experience greater volatility in response to adversity – we conducted two standard statistical tests: a Welch’s t-test to compare means and a Levene’s test to compare variances.

We began by calculating the mean change in IPM for all players following a midseason injury to a high average IPM teammate, grouped by whether their team was classified as spiky or uniform. Contrary to our initial expectations, the Welch’s t-test revealed no statistically significant difference in average IPM change between the two team types. In other words, on average, players in both spiky and uniform teams adjusted similarly in terms of performance magnitude.

However, our theoretical model did not rest on mean outcomes, it rested on the idea that spiky teams lack structural redundancy, and therefore would respond to disruption with greater variability. To test this, we applied a Levene’s test for homogeneity of variance to the same two groups. This time, the results were conclusive: the variance in IPM change was significantly higher among spiky teams ($\sigma^2 \approx 21.3$) than among uniform teams ($\sigma^2 \approx 6.9$), with a p-value less than 0.0001.

This result directly supports our hypothesis. While the average level of chemistry change may be comparable across teams, the stability of those changes is not. Uniform teams responded to adversity in a more predictable and consistent manner, while spiky teams saw disproportionate fluctuations, some players improved dramatically, others collapsed. This reflects a lack of structural balance: when chemistry is concentrated among a few dominant links, any disturbance to those links sends unpredictable ripple effects across the network.

Our takeaway is that team resilience under pressure is not defined by how much chemistry changes on average, but by how predictably it changes across players. This has critical implications for roster construction, game planning, and trade decisions. Teams built on uniform chemistry structures can absorb disruption and reallocate responsibility smoothly. Spiky teams, by contrast, face a roll of the dice, sometimes rebalancing effectively, sometimes falling apart.

10. Conclusion and Drawbacks

This study set out to investigate whether the structure of a team’s chemistry network – specifically, the distribution of effective player pairings – could explain how resilient a team is to adversity in the form of a midseason injury. We hypothesized that teams with uniformly distributed chemistry would show more consistent performance post-injury, while spiky teams, those relying on a few elite connections, would exhibit more volatile adjustments.

Our findings support this hypothesis, but with an important nuance. A Welch's t-test found no significant difference in the average change in chemistry between the two groups. However, the Levene's test revealed a striking difference in variance: chemistry change was significantly more volatile in spiky teams ($\sigma^2 \approx 21.3$) than in uniform ones ($\sigma^2 \approx 6.9$), with a p-value < 0.0001 . This suggests that team resilience is not about whether chemistry shifts, but about how predictably it shifts. Uniform teams absorb disruption with stability; spiky teams react chaotically, sometimes adapting well, sometimes collapsing.

The implication this may hold for team construction are:

- Playoff-caliber teams may benefit from targeting role players who strengthen redundant chemistry links, thus enhancing resilience.
- Rebuilding teams, on the other hand, may intentionally embrace spikiness and volatility, gambling on a single star's transformative potential.

That said, there are important limitations to our approach:

Limitations of the Chemistry Metric (IPM):

IPM is built from box-score statistics shared during court time. While it captures many on-ball outcomes, assists, points, rebounds, etc., it misses off-ball movement, leadership, rotations, and communication. True chemistry encompasses more than shared stats gained during shared playtime.

Interpretation of Betweenness Centrality:

In theory, betweenness should help identify players who glue teams together. But because our edges are based on IPM, betweenness often highlights players with many minutes across rotations, not necessarily the star-based structural connections we deemed important for analysis.

Scope of Data:

We focused on four representative teams. A broader sample would allow for more robust regression models, particularly those that correlate IPM volatility with win-rate changes, providing a stronger link between chemistry and outcomes people care about.

Simulated vs. Real-World Removal:

Our model assumes injuries are disruptive enough to simulate structural holes. In practice, coaching adjustments, strength of schedule, and opponent quality also affect post-injury outcomes—factors that we could not fully control for.

Despite these constraints, our work introduces a foundational framework for thinking about team chemistry: not just as “who plays well together,” but as a networked system of redundancies. Resilience lies not in any single connection, but in a structure that can adapt without breaking.

Sources (Literature Review):

Mukherjee, S., Huang, Y., Neidhardt, J., Uzzi, B., & Contractor, N. (2018). Prior shared success predicts victory in team competitions. *Nature Human Behaviour*, 3(1), 74–81.
<https://doi.org/10.1038/s41562-018-0460-y>

Pradeep, N. (2024). Identifying correlations between NBA success and team dynamics: A statistical analysis. *The National High School Journal of Science*, 1–7.
<https://nhsjs.com/wp-content/uploads/2024/06/Identifying-Correlations-Between-NBA-Success-and-Team-Dynamics-A-Statistical-Analysis.pdf>

Fewell, J. H., Armbruster, D., Ingraham, J., Petersen, A., & Waters, J. S. (2012). Basketball teams as strategic networks. *PLoS ONE*, 7(11), e47445.
<https://doi.org/10.1371/journal.pone.0047445>