

Stability and Fairness in Strategic Systems: Nash Equilibrium and the Shapley Value

Nino Miljkovic (January, 2026)

Table of Contents

1. Introduction
2. Part I: Nash Equilibrium – Seeking Stability
3. Part II: Shapley Value - Fairness Through Marginal Contribution
4. Part III: Nash and Shapley Together - Design, Not Description
5. Conclusion

Introduction

Game theory exists because intuition, while powerful, systematically fails in strategic environments. When outcomes depend not only on what we choose but also on how others respond, individual reasoning alone is insufficient [1]. The field emerged from the recognition that traditional economic and mathematical tools were inadequate for analyzing situations involving strategic interdependence, contexts where each agent's optimal choice depends critically on the choices of others. Two of the most influential concepts in game theory, Nash equilibrium and the Shapley value, address different but deeply connected problems that arise in such environments.

Nash equilibrium answers the question of stability: *Given a set of incentives and rules, what patterns of behavior will persist once everyone has adapted?* First formalized by John Nash in his groundbreaking 1950 paper [2], this concept revolutionized economics, political science, and biology by providing a rigorous framework for predicting strategic behavior. The significance of this contribution was recognized with the Nobel Memorial Prize in Economic Sciences in 1994.

The Shapley value answers the question of fairness and attribution: *When multiple actors together create value, how should that value be divided among them?* Introduced by Lloyd Shapley in 1953 [3], this solution concept provides a unique, axiomatically justified method for allocating payoffs in cooperative games. Shapley's contribution, recognized with his own Nobel Prize in 2012, has proven remarkably versatile, finding applications far beyond its original mathematical context.

These concepts are often introduced separately in textbooks and courses, but, in my opinion, their real power emerges when they are understood together. Nash equilibrium tells us what will happen if incentives are left unchanged; Shapley value gives us a principled way to redesign incentives so that better equilibria become possible. This complementarity has become increasingly important in modern applications, from mechanism design [4] to algorithmic fairness [5] to understanding complex socio-technical systems. This essay explores both ideas in depth, explains why they matter beyond their formal definitions, and shows how they apply to modern domains such as business strategy, machine learning, and organizational design.

Part I: Nash Equilibrium – Seeking Stability

1.1 The Core Question Nash Answers

At its heart, Nash equilibrium answers a simple but profound question: *Given what everyone else is doing, do I have an incentive to change my behavior?* [2]

If the answer is "no" for every participant, the system is in a Nash equilibrium. Formally, a strategy profile $(s_1^*, s_2^*, \dots, s_n^*)$ is a Nash equilibrium if for every player i , their strategy s_i^* is

a best response to the strategies of all other players. Mathematically, this means that for all players i and all alternative strategies s_i :

$$u_i(s_i^*, s_{-i}^*) \geq u_i(s_i, s_{-i}^*)$$

where u_i represents player i 's utility function and s_{-i}^* denotes the strategies of all players except i .

Importantly, this does not mean the outcome is efficient, fair, or desirable. It only means that no single actor can improve their situation by changing strategy alone, a property known as *individual rationality* under fixed expectations. This distinction between *stability* and *quality* is where intuition often fails. People tend to assume that bad outcomes require bad actors or irrational behavior. Nash equilibrium shows that even perfectly rational individuals, each acting in their own best interest with full information, can collectively produce outcomes that everyone dislikes. This is a phenomenon that extends the classical notion of the "invisible hand" to show that individual rationality does not guarantee collective optimality.

1.2 A Simple Illustration: The Prisoner's Dilemma and Price Competition

Consider a classic competitive scenario: two firms deciding whether to engage in a price war or keep prices high. This situation can be modeled as a Prisoner's Dilemma, one of the most studied games in strategic analysis. Each firm reasons as follows:

- If the other firm keeps prices high, undercutting yields higher market share and short-term profits
- If the other firm cuts prices, matching the cut prevents catastrophic market share loss

The payoff structure creates a dominant strategy: regardless of what the competitor does, cutting prices yields a better individual outcome. The result is mutual price cutting, even though both firms would prefer the higher-price outcome. In the language of game theory, "High Price, High Price" is Pareto efficient, referring to the fact that no player can be made better off without making another worse off, but it is not a Nash equilibrium because each firm has an incentive to deviate. Once both firms are cutting prices, neither can improve profits by unilaterally raising prices. This outcome is stable, hence a Nash equilibrium, but clearly suboptimal from the collective perspective.

Real-world examples abound. The airline industry has repeatedly demonstrated this dynamic, with price wars driving several carriers to bankruptcy despite high consumer demand. Similarly, the "race to the bottom" in online advertising rates reflects an equilibrium where platforms compete away margins despite providing valuable services. The telecommunications sector has experienced similar patterns, where infrastructure investment games lead to either over-investment (when competitors build redundant networks) or under-investment (when each waits for others to invest first).

The key lesson is that Nash equilibrium describes what behavior settles into, not what behavior should occur. This distinction becomes crucial in policy design, regulation, and organizational management.

1.3 Why Nash Equilibrium Is Powerful

Nash equilibrium matters because it explains why certain problems persist despite intelligence, goodwill, and experience. Its explanatory and predictive power stems from several key properties:

1. **Self-enforcing outcomes**

Some outcomes persist precisely because any individual deviation is punished by the system. Appeals to ethics, cooperation, or "better culture" fail when incentives remain unchanged. This explains why corporate codes of conduct often prove ineffective without accompanying changes to compensation structures, promotion criteria, or performance metrics. The concept of "self-enforcing agreements" in contract theory builds directly on Nash's insight that sustainability requires individual incentive compatibility.

In environmental economics, this principle explains the difficulty of achieving voluntary emissions reductions. Even when all parties recognize the collective benefit of reduced pollution, individual countries or firms face strong incentives to free-ride on others' efforts [6]. The Paris Agreement's architecture attempts to address this by creating reputational costs and verification mechanisms that alter the equilibrium incentives.

2. **Predictive power**

Nash equilibrium allows us to predict behavior *before* observing it. If a proposed policy creates a dominant incentive to defect, widespread defection is not a moral failure, it is an equilibrium. This predictive capability has proven invaluable in auction design, where understanding equilibrium bidding strategies enables designers to choose formats that maximize revenue or efficiency. The FCC spectrum auctions, which have allocated hundreds of billions of dollars in telecommunications rights, rely heavily on Nash equilibrium analysis to predict bidder behavior and prevent collusion.

In computational settings, the predictive power of Nash equilibrium has been formalized through complexity theory. The PPAD (Polynomial Parity Arguments on Directed graphs) complexity class, which characterizes the difficulty of computing Nash equilibria, has revealed deep connections between game theory and computational complexity [7]. This work shows that while Nash equilibria are guaranteed to exist in finite games [2], finding them can be computationally intractable. This result has profound implications for algorithmic game theory and market design.

3. **Design implications**

Once an equilibrium is identified, the only way to change behavior is to change payoffs, constraints, or available strategies. This insight underpins mechanism

design, the "engineering" side of game theory that seeks to design institutions and rules that produce desired equilibria [4]. The 2007 Nobel Prize awarded to Leonid Hurwicz, Eric Maskin, and Roger Myerson recognized this paradigm's importance in economics and beyond.

Modern platform design heavily leverages these principles. Ride-sharing companies like Uber and Lyft use dynamic pricing to maintain equilibrium between driver supply and passenger demand. Social media platforms design algorithmic feeds and moderation policies to shift user behavior toward desired equilibria, though these efforts often produce unintended consequences when the full strategic landscape is not considered.

1.4 Nash Equilibrium in Dynamic and Real Systems

In real systems, Nash equilibria are rarely static. They evolve as participants learn, adapt, and respond to feedback, which is a phenomenon studied in evolutionary game theory and learning in games [8]. The concept of Evolutionary Stable Strategy (ESS), introduced by Maynard Smith and Price [9], extends Nash equilibrium to biological contexts where strategies evolve through natural selection rather than rational choice.

- In sports, rule changes (such as introducing the three-point line in basketball in 1979) reshape equilibria by altering payoffs, forcing teams to adopt new strategies. The three-point line fundamentally changed optimal team composition and playing style, creating new equilibria that emphasize perimeter shooting and floor spacing. Analytics have revealed that the value of three-point shooting was systematically undervalued for decades, and teams that recognized this gained significant competitive advantages. Similarly, defensive strategies evolved in response, with zone defenses and switching schemes becoming more prevalent as equilibrium strategies adapted to the new offensive landscape.
- In business, competitive markets converge toward pricing, advertising, and investment patterns that no firm can easily escape. The software industry provides a striking example: the shift to freemium models and subscription pricing represents an equilibrium response to near-zero marginal costs and network effects. Once dominant firms adopt these models, competitors face strong pressure to follow, even when alternative pricing structures might generate higher industry-wide profits. The emergence of open-source software similarly reflects an equilibrium where firms contribute to public goods because proprietary alternatives have become uncompetitive.

Financial markets exhibit dynamic equilibria where trading strategies, asset prices, and risk management practices co-evolve. The 2008 financial crisis revealed how a stable-seeming equilibrium (the "Great Moderation") could conceal growing systemic risks, and how the system could rapidly transition to a new, far worse equilibrium once key assumptions broke down. Modern financial regulation increasingly focuses on understanding these equilibrium dynamics rather than simply constraining individual behaviors.

- In machine learning systems, users and adversaries adapt to models, forming new equilibria that may degrade performance over time. We refer to this phenomenon as performative prediction [10]. Credit scoring models change borrower behavior, which in turn affects the distribution of data the model encounters in deployment. Spam filters drive spammers to adopt new evasion techniques, producing an arms race that resembles biological co-evolution. Recommender systems shape user preferences even as they attempt to predict them, creating feedback loops that can reinforce existing patterns or drive systems toward new equilibria [11].

A deployed ML model is not merely making predictions; it is selecting an equilibrium by shaping incentives and behavior. This recognition has spawned new research areas at the intersection of machine learning and game theory, including strategic classification [12], performative prediction [10], and algorithmic mechanism design. These fields recognize that models operate in strategic environments where rational agents respond to the model's predictions and decisions.

1.5 The Dark Side of Nash Equilibrium

The most uncomfortable implication of Nash equilibrium is this: Some bad outcomes are not accidents. They are inevitable under the current rules.

This explains several persistent pathologies in modern systems:

- **Metric gaming in organizations:** When performance is measured by narrow metrics, employees rationally optimize for those metrics even when doing so harms broader organizational goals [13]. Teachers "teaching to the test," doctors focusing on billable procedures over patient outcomes, and software engineers optimizing for lines of code rather than functionality all reflect equilibrium responses to misaligned incentives. Campbell's Law formalizes this observation: "The more any quantitative social indicator is used for social decision-making, the more subject it will be to corruption pressures and the more apt it will be to distort and corrupt the social processes it is intended to monitor" [14].
- **Over-optimization in machine learning:** Models trained on historical data can learn to exploit spurious correlations and dataset artifacts, achieving high benchmark performance while failing catastrophically in deployment [15]. Adversarial examples, such as inputs specifically designed to fool neural networks, reveal that standard training procedures converge to equilibria that are brittle and exploitable [16]. The phenomenon of "shortcut learning" shows that models will rationally exploit whatever features provide the strongest training signal, even when those features are unreliable or unrelated to the true underlying task [15].
- **Arms races in business and geopolitics:** Competitive dynamics can drive all parties to invest heavily in zero-sum competition, dissipating resources that could create value elsewhere. The Cold War arms race represents the canonical example, where both superpowers accumulated nuclear arsenals far beyond any conceivable

military need, each responding rationally to the other's buildup. In business, advertising wars, patent thickets, and excessive investment in low-margin products reflect similar dynamics. The pharmaceutical industry's focus on "me-too" drugs rather than breakthrough innovations partially reflects an equilibrium where copying successful drugs carries less risk than fundamental research.

- **Toxic equilibria in online platforms:** Social media platforms can converge to equilibria where outrage, misinformation, and polarization dominate because these contents generate engagement and ad revenue [17]. Content moderation policies face the challenge that any rule will be gamed by strategic actors seeking to stay just within the boundaries of acceptable behavior. The result is often a race to the bottom where platforms compete on addictiveness rather than user welfare. Recent research in computational social science has begun to model these dynamics formally, revealing how platform design choices affect equilibrium content distributions and user behavior [18].

Recognizing a Nash equilibrium prevents wasted effort on persuasion and instead focuses attention on redesigning the system itself. As mechanism design has taught us, changing outcomes requires changing the game, not merely exhorting players to make different choices [4].

Part II: Shapley Value - Fairness Through Marginal Contribution

2.1 The Problem Shapley Answers

While Nash equilibrium focuses on strategic stability, the Shapley value addresses a different question: *When a group of participants collectively creates value, how should that value be divided among them?* [3]

This problem arises whenever outcomes depend on cooperation, a situation formalized in cooperative game theory. Unlike non-cooperative games where the Nash equilibrium concept applies, cooperative games assume that binding agreements are possible and focus on how coalitions form and divide payoffs. The challenge is that value creation is often superadditive, meaning that the whole exceeds the sum of the parts, making individual contributions difficult to disentangle.

This problem appears across diverse domains:

- **Teams delivering a project:** Software development teams produce complex systems where individual contributions are deeply interdependent. A brilliant algorithm is useless without robust engineering, effective testing, and thoughtful user interface design. How should credit be allocated for project success?

- **Channels contributing to a sale:** In multi-touch attribution, customers interact with numerous marketing touchpoints before purchasing. Email campaigns, social media ads, search engine marketing, and direct sales contacts all play roles, but which deserve credit? [19] Traditional "last-click" attribution ignores the contribution of earlier touchpoints that built awareness and interest.
- **Features contributing to a model prediction:** In machine learning, predictions emerge from complex interactions among hundreds or thousands of features. Linear models with interaction terms, tree ensembles, and neural networks all combine features in non-linear ways that obscure individual contributions. For high-stakes applications like loan decisions or medical diagnosis, understanding feature contributions is essential for transparency and fairness [20].

“Naïve” approaches, such as equal splitting, seniority, or visibility, often fail because they ignore marginal contribution. Equal splitting rewards free-riding and ignores differences in actual impact. Seniority-based systems reward tenure over performance and discourage innovation. Visibility-based allocation favors public-facing work over essential but less visible contributions; a dynamic that systematically undervalues the "invisible work" that keeps organizations functioning.

2.2 The Core Idea of the Shapley Value

The Shapley value assigns credit to each participant based on a simple principle: A participant's share should equal their average marginal contribution across all possible ways the group could have formed [3].

Formally, for a cooperative game with player set N and characteristic function v (which assigns a value to each coalition $S \subseteq N$), the Shapley value $\varphi_i(v)$ for player i is:

$$\varphi_i(v) = \sum_{S \subseteq N \setminus \{i\}} [|S|! (|N| - |S| - 1)! / |N|!] \times [v(S \cup \{i\}) - v(S)]$$

This formula can be interpreted as follows:

1. Consider all possible orders in which players might join a coalition (there are $|N|!$ such orderings)
2. For each ordering, calculate how much value player i adds when they join
3. Average these marginal contributions across all orderings

The weighting factor $[|S|!(|N|-|S|-1)!]/|N|!$ represents the probability that the players in S arrive before player i , and the players in $N \setminus (S \cup \{i\})$ arrive after i , when players join in a uniformly random order.

This process ensures that credit reflects both standalone value and synergy with others. The Shapley value is the unique allocation method that satisfies four intuitively appealing axioms [3]:

1. **Efficiency:** The sum of all players' Shapley values equals the total value created by the grand coalition
2. **Symmetry:** If two players contribute identically to all coalitions, they receive equal shares
3. **Dummy player:** If a player contributes nothing to any coalition, their share is zero
4. **Additivity:** For two games v and w , the Shapley value of their sum is the sum of their individual Shapley values

These axioms provide an axiomatic foundation for fairness that has proven remarkably robust across applications. The uniqueness theorem, stating that these four axioms uniquely determine the Shapley value, gives it normative force: if we accept these principles, we must accept Shapley value as the fair allocation [3].

2.3 Why This Is Not Obvious

Human intuition tends to overvalue several factors that the Shapley value corrects:

- **Visibility:** Work that is public-facing or highly visible often receives disproportionate credit. The salesperson who closes a deal may receive full credit while the product developers, customer support agents, and marketing teams who enabled the sale are overlooked. Research on collaborative scientific work shows that last authors on papers (typically senior researchers) often receive disproportionate credit relative to their actual contributions.
- **Final actions:** The person who completes a project or makes the final decision often appears most important, even when earlier contributions were essential. In sports, the player who scores the winning goal receives accolades, while the defenders, midfielders, and goalkeeper who made the victory possible remain less celebrated. This "recency bias" systematically distorts attribution.
- **Power or negotiation leverage:** Those with bargaining power or political influence often capture more value than their contributions justify. While the Nash bargaining

solution and other cooperative solution concepts account for outside options and bargaining power, the Shapley value deliberately ignores these factors to focus purely on contribution. This makes it valuable as a baseline for assessing whether actual allocations reflect merit or merely power dynamics.

The Shapley value corrects these biases by recognizing **enablers**: participants who may contribute little alone but unlock large value in combination with others. A specialist who provides a crucial technical insight might seem less important than a generalist who coordinates the team, but the Shapley value captures that without the specialist's contribution, the entire project might fail. This makes it especially powerful in complex systems where contributions are interdependent.

2.4 Shapley Value in Practice

The Shapley value has found applications in an remarkably diverse array of domains, often producing surprising insights that challenge conventional wisdom.

Business and Economics

- **Revenue attribution across marketing channels:** Multi-touch attribution using Shapley values reveals which marketing channels truly drive conversions versus which merely intercept customers already planning to purchase [19, 21]. Companies like Google and Facebook have developed Shapley-based attribution models to help advertisers understand channel effectiveness. Research shows that Shapley attribution often dramatically differs from simple last-click or first-click models, redistributing credit to awareness-building activities that traditional models undervalue [19].
- **Profit sharing in joint ventures:** When multiple firms collaborate on a project, Shapley value provides a principled basis for dividing profits. The European Union has used Shapley-value-based approaches to allocate costs and benefits in transnational infrastructure projects. In pharmaceutical development, where multiple companies may contribute different capabilities (research, clinical trials, manufacturing, distribution), Shapley value offers a fair baseline for profit-sharing agreements.
- **Bonus allocation in collaborative teams:** Rather than dividing bonuses equally or based on seniority, Shapley-based approaches reward actual contribution to team outcomes. Several technology companies have experimented with Shapley-inspired performance evaluation systems, though implementation challenges remain significant. The key advantage is that Shapley value naturally accounts for

complementarities: a highly specialized expert might have low standalone value but high Shapley value because they enable other team members to be effective.

Sports Analytics

Sports analytics has embraced Shapley value as a tool for measuring player contributions that traditional statistics miss:

- **Valuing defensive players whose impact is indirect:** In basketball, defensive specialists who prevent opposing players from receiving the ball or force them into low-percentage shots often have minimal traditional statistics (blocks, steals, rebounds) yet contribute significantly to winning. Shapley-based metrics like Defensive Real Plus-Minus attempt to capture these contributions by measuring how team defensive performance changes with and without specific players [22].
- **Measuring off-ball contributions:** In soccer, players who create space for teammates, draw defenders away from dangerous areas, or provide passing options rarely show up in assist statistics. Shapley value analysis of passing networks reveals which players most effectively enable team attacking play, even when they don't directly participate in goals [23].
- **Explaining why certain role players are indispensable:** Teams often feature "glue guys"—players with unremarkable individual statistics who nevertheless prove essential to team success. Shapley value analysis can reveal that these players have high value because they complement star players, enabling the stars to perform optimally. This insight has influenced roster construction and salary negotiations, as teams recognize that traditional statistics undervalue complementary skills.

Machine Learning: SHAP and Explainable AI

The Shapley value underpins SHAP (SHapley Additive exPlanations), one of the most widely used explainability techniques in modern ML [24]. SHAP has become a cornerstone of the explainable AI movement, bridging game theory and machine learning in a powerful synthesis.

In the SHAP framework:

- Players become features
- Value becomes a model's prediction

- Shapley values quantify each feature's contribution to moving the prediction away from a baseline (typically the average prediction)

For a prediction $f(x)$ where x represents a feature vector, SHAP values satisfy:

$$f(x) = \varphi_0 + \sum_i \varphi_i$$

where φ_0 is the base value (expected model output) and φ_i is the SHAP value of feature i [24].

This decomposition enables:

1. **Transparency:** Stakeholders can understand why a model made a particular prediction. In criminal justice, where algorithms increasingly inform bail and sentencing decisions, SHAP values help judges understand the factors driving risk assessments. In lending, SHAP explanations help satisfy adverse action notice requirements that mandate explanations for loan denials [25].
2. **Fairness audits:** By examining how sensitive attributes (race, gender, age) contribute to predictions, researchers can detect and quantify discrimination [5]. SHAP values reveal not just whether a model uses protected attributes, but how strongly and in what combinations. This has proven crucial for identifying proxy discrimination, where models use seemingly neutral features that correlate with protected attributes.
3. **Debugging:** When models perform unexpectedly, SHAP values help identify whether the model is exploiting spurious correlations or dataset artifacts [26]. In medical imaging, SHAP analysis revealed that some models were identifying X-ray machines rather than pathology. This poses as a dangerous shortcut that traditional accuracy metrics missed [27]. In NLP, SHAP has exposed models relying on annotation artifacts and dataset biases [28].
4. **Regulatory compliance:** As AI regulation emerges globally, model explainability is becoming mandatory for high-stakes applications. The EU's General Data Protection Regulation (GDPR) includes a "right to explanation" for automated decisions, and proposed AI regulations explicitly require explainability for high-risk systems [29]. SHAP provides a mathematically principled, legally defensible explanation method.

However, computing exact SHAP values is computationally expensive—requiring evaluating the model on all 2^n possible feature subsets [24]. This has spawned research on efficient approximations, including TreeSHAP for tree-based models [30], KernelSHAP for model-agnostic explanations [24], and DeepSHAP for neural networks. Recent work has

also investigated the statistical properties of SHAP value estimates and developed methods for quantifying their uncertainty [31, 32].

Part III: Nash and Shapley Together - Design, Not Description

3.1 Complementary Roles

Nash equilibrium and the Shapley value address different layers of the same system, and their interaction reveals deep insights about mechanism design and organizational effectiveness:

- **Nash equilibrium** describes *behavioral stability*: Given a set of rules and incentives, what patterns of behavior will emerge and persist? [2]
- **Shapley value** describes *value attribution*: Given observed outcomes, how should credit or payoffs be allocated among participants? [3]

One tells us what will happen; the other helps us decide how outcomes should be measured and rewarded. This complementarity has profound implications for system design.

Traditional economic analysis often treats institutions as fixed and studies the equilibria they produce. Mechanism design reverses this logic: specify the desired equilibrium, then design institutions that make it stable [4]. The Shapley value provides a principled target for these designs.

3.2 How Shapley Can Reshape Nash Equilibria

This connection is crucial for practical applications. Incentives shape equilibria, and Shapley-based attribution can be used to redesign incentives toward more desirable equilibria.

Consider an ML organization developing a complex product. Different teams contribute to various components:

- Data engineering builds pipelines and ensures data quality
- Feature engineering creates predictive signals
- Model development builds and tunes algorithms
- Infrastructure provides computational resources and deployment systems

- Product management defines requirements and prioritizes work

Without proper attribution, teams optimize local metrics, leading to a Nash equilibrium of metric gaming and siloed optimization [13]. The data team optimizes for pipeline throughput regardless of data quality. Feature engineers create numerous low-value features to appear productive. Model developers chase benchmark performance using features that won't reliably deploy. Infrastructure over-invests in computational capacity to avoid blame for resource constraints. Product managers proliferate requirements to demonstrate engagement.

Each team's behavior is individually rational given the incentives, but collectively the organization produces brittle, unreliable systems that fail in deployment [33].

By introducing Shapley-based credit allocation for shared outcomes—overall model performance, business metrics, user satisfaction—we can align rewards with true contribution. This changes incentives, thus producing a new, more cooperative equilibrium where:

- Data engineers optimize for data quality because it increases their Shapley value
- Feature engineers focus on genuinely predictive features that improve end-to-end performance
- Model developers consider deployment constraints because they affect realized value
- Infrastructure invests optimally because their Shapley value reflects enabling others' success
- Product managers prioritize high-impact requirements because success metrics matter more than activity metrics

In this sense, Shapley values are not merely descriptive, they are mechanism design tools. They transform the game from a non-cooperative equilibrium of local optimization to a cooperative equilibrium of aligned interests.

Real-world implementations face challenges. Computing Shapley values for complex organizations requires defining appropriate value functions, handling uncertainty in contributions, and maintaining computational tractability. Several companies have developed practical approximations:

- LinkedIn uses Shapley-inspired attribution for feed ranking improvements
- Microsoft applies Shapley value concepts to evaluate A/B test contributions

- Various startups offer Shapley-based analytics for marketing attribution [19]

Research continues on overcoming implementation barriers, including efficient computation [34], handling strategic manipulation, and incorporating dynamic considerations.

3.3 Stability vs Fairness Tension

An important insight is that Nash equilibria can be unfair (producing inequality unrelated to contribution), and fair allocations may not be stable unless incentives support them. This tension appears throughout economics and political philosophy:

- **The bargaining problem:** Nash himself studied how players divide surplus through bargaining [35]. While the Nash bargaining solution accounts for both fairness and stability considerations, it often differs from the Shapley value because it incorporates outside options and bargaining power. When powerful players can credibly threaten to walk away, stable agreements may require giving them more than their Shapley value; an inequality reflecting power rather than contribution.
- **Constitutional design:** Democratic institutions face the challenge of creating voting rules that are both stable (encouraging compliance rather than revolution) and fair (respecting equality and minority rights). Shapley-Shubik power indices measure voting power in weighted voting systems, revealing when formal voting rules diverge from true influence [36]. This analysis has influenced institutional design in the European Union, the United Nations Security Council, and corporate governance.
- **Coalition formation:** In parliamentary systems, government coalitions must be both stable (parties won't defect) and roughly fair (coalition partners receive ministerial positions commensurate with their contribution to the majority). Shapley value analysis predicts which coalitions will form and how cabinet positions are allocated.

Therefore, effective system design must address both:

1. **What behaviors will persist?** (Nash equilibrium analysis)
2. **How should rewards be distributed to support those behaviors?** (Shapley value as a fairness benchmark)

Ignoring the first question leads to proposals that sound good but prove unstable in practice. Ignoring the second question produces stable but unjust systems that waste

human potential and generate legitimate grievances. The intersection of game theory and political philosophy increasingly recognizes this dual requirement.

3.4 Implications for Modern Systems

Machine Learning Systems

ML systems increasingly interact with strategic human agents, creating feedback loops that traditional supervised learning ignores [37]. Once deployed, they:

- **Influence behavior:** Credit scoring models change borrowing behavior; content recommendation shapes user preferences [11]; facial recognition systems induce strategic adaptation in presentation [38]
- **Create feedback loops:** Predictive policing concentrates enforcement in predicted high-crime areas, which generates more data from those areas, reinforcing initial patterns regardless of true crime distribution [39]; hiring algorithms trained on biased historical data perpetuate and amplify existing inequities [40]
- **Invite strategic adaptation:** Adversarial attacks on image classifiers [16]; resume optimization to exploit hiring algorithms [12]; manipulation of recommendation systems

Nash equilibrium explains why systems drift, get gamed, or collapse [10]. Models that ignore strategic responses often fail dramatically when deployed. The emerging field of performative prediction formalizes these dynamics, analyzing how predictions change the distribution they aim to predict [10]. Strategic classification studies how machine learning performs when individuals strategically modify their features to receive favorable predictions [12].

Shapley value explains how to fairly attribute responsibility, detect bias, and align incentives across teams and features [24, 5]. When models make consequential decisions, Shapley-based explanations help identify whether outcomes reflect legitimate predictive factors or problematic proxies and spurious correlations [27, 28]. This proves essential for accountability: if an automated system causes harm, Shapley value analysis can help determine which components or design choices bear responsibility.

The combination of Nash equilibrium thinking and Shapley value tools enables better ML system design:

- Anticipate strategic responses using game-theoretic modeling [41]
- Design robust systems that maintain performance under adaptation [42]

- Allocate credit appropriately across teams using SHAP for model components
- Create incentive-compatible systems where honest reporting is an equilibrium [43]

Business and Organizations

Organizations are not machines executing predetermined algorithms; they are strategic systems where individuals and teams respond to incentives. Nash equilibrium explains why certain dysfunctional behaviors persist despite leadership's intentions and employees' goodwill:

- **Silo formation:** When teams are rewarded for local metrics, cross-functional collaboration becomes a Nash equilibrium defection. Even when collaboration would benefit the organization, no individual team can afford to sacrifice their metrics for collective gain.
- **Innovation resistance:** Established business units may rationally resist disruptive innovations that cannibalize existing revenue, even when innovation is essential for long-term survival. This is not short-sightedness but equilibrium behavior given typical incentive structures.
- **Information hoarding:** When knowledge is power, sharing information can reduce individual value. Absent explicit incentives for knowledge sharing, information silos emerge as a Nash equilibrium.
- **Risk aversion:** When failure is punished more severely than success is rewarded, employees rationally avoid ambitious projects. This risk-averse equilibrium stifles innovation even in organizations that claim to value entrepreneurship.

Shapley value provides a principled way to allocate credit, bonuses, and responsibility in collaborative environments. By measuring true contribution rather than visibility or power, Shapley-based systems can shift equilibria toward productive collaboration:

- Reward cross-functional collaboration by measuring impact on shared outcomes
- Credit enablers who make others successful, not just individual contributors
- Recognize indirect contributions that traditional metrics miss
- Create incentives for knowledge sharing and mentorship

Several forward-thinking organizations have experimented with these approaches. W.L. Gore & Associates, famous for its flat organizational structure, uses peer-based evaluation systems that approximate Shapley value thinking by assessing individuals' contributions to

collective outcomes. Research continues on making these approaches more systematic and scalable.

Governance and Policy

Regulation and policy are fundamentally about changing equilibria through shifting behavior from undesirable to desirable patterns. Understanding Nash equilibria prevents naive policy design that ignores strategic responses:

- **Tax policy:** Tax rates affect not just compliance but economic behavior, such as work effort, investment, corporate structure, and international capital flows. Optimal tax design must account for these equilibrium responses. The Laffer curve illustrates that higher tax rates don't necessarily yield higher revenue once behavioral responses are considered. Modern tax policy increasingly uses mechanism design principles to create incentives for honest reporting while maintaining revenue.
- **Environmental regulation:** Cap-and-trade systems for carbon emissions work by creating a market where the Nash equilibrium involves efficient pollution reduction [44]. The European Union Emissions Trading System demonstrates both the promise and challenges of this approach. When permits are allocated suboptimally or monitoring is inadequate, strategic behavior can undermine effectiveness.
- **Financial regulation:** Post-2008 reforms aim to change equilibria in financial markets by altering capital requirements, limiting certain activities, and improving transparency. However, regulatory arbitrage where firms strategically reorganize to minimize regulatory burden represents a Nash equilibrium response that regulators must anticipate. The shadow banking system emerged partially as an equilibrium response to traditional bank regulation.
- **Platform governance:** Online platforms must design rules anticipating strategic responses from users, advertisers, and content creators [45]. Content moderation policies face the challenge that any rule creates incentives to stay just within boundaries. Effective platform governance requires understanding these equilibrium dynamics.

Shapley-style attribution informs fair taxation, liability assignment, and compensation in policy contexts:

- **Tax burden allocation:** Shapley value can determine how to fairly distribute tax obligations across jurisdictions for multinational corporations, accounting for where value is truly created. This has become crucial as digital businesses generate revenue across borders with minimal physical presence.

- **Climate responsibility:** Attributing responsibility for climate change and determining fair contributions to mitigation efforts involves Shapley-value-like reasoning about historical emissions, current capacity, and future obligations [46]. The principle of "common but differentiated responsibilities" in climate agreements reflects this logic.
- **Liability in complex systems:** When harm results from multiple parties' actions (pharmaceutical companies, doctors, insurers in a medical error; developers, manufacturers, operators in an autonomous vehicle accident) Shapley value provides a framework for attributing responsibility proportional to causal contribution.
- **Public goods provision:** When multiple jurisdictions benefit from public infrastructure or research, Shapley value can determine cost-sharing arrangements. The European Union uses such approaches for transnational projects like transportation networks and research initiatives.

3.5 Bridging Theory and Practice

The gap between theoretical elegance and practical implementation remains significant. Several challenges confront practitioners attempting to apply these concepts:

1. **Computational complexity:** Computing Nash equilibria is PPAD-complete [7]; computing Shapley values requires exponential calculations [3]. Practical applications require approximations, heuristics, or restrictions to tractable special cases.
2. **Information requirements:** Both concepts assume knowledge often unavailable in practice. Nash equilibrium requires knowing others' payoffs; Shapley value requires the value function for all coalitions. Incomplete information and uncertainty complicate real-world applications.
3. **Multiplicity and selection:** Many games have multiple Nash equilibria, raising the equilibrium selection problem [47]. Without additional principles for choosing among equilibria, predictive power diminishes. Similarly, different value functions yield different Shapley values, and defining appropriate value functions for complex organizations is non-trivial.
4. **Dynamic considerations:** Both concepts are primarily static, but real systems evolve over time. While dynamic extensions like repeated games, stochastic games, and dynamic cooperative games exist, they add substantial complexity [48].

5. **Behavioral realism:** Both concepts assume rational optimization, but humans exhibit systematic deviations from rationality [49]. Behavioral game theory studies these deviations and their implications [50]. Yet the insights from classical game theory often remain valuable even when rationality is bounded.

Despite these challenges, the conceptual frameworks remain invaluable. They provide:

- A language for reasoning about strategic interaction
- Benchmarks against which real behavior can be measured
- Design principles for creating better institutions [4]
- Diagnostic tools for understanding system dysfunction

As computational methods improve and data becomes more abundant, practical applications expand. Machine learning itself provides tools for overcoming some limitations: learning equilibria from data, approximating Shapley values efficiently [32], and handling high-dimensional strategy spaces.

Conclusion

Nash equilibrium and the Shapley value represent two pillars of strategic reasoning, each addressing fundamental questions about how systems behave and how outcomes should be distributed.

Nash equilibrium teaches us humility: even intelligent, well-intentioned actors can become trapped in undesirable outcomes when incentives are misaligned [2]. It forces us to confront the difference between what people *want* to do and what they *can afford* to do within a system. This insight has revolutionized economics, transformed political science, reshaped biology through evolutionary game theory [9], and increasingly influences computer science [7] and machine learning [10]. The concept's power lies not in guaranteeing optimal outcomes but in explaining why suboptimal outcomes persist and what must change to improve them.

The Shapley value teaches us fairness with rigor: attribution should be based not on power, visibility, or narrative, but on marginal contribution averaged across contexts [3]. It provides a principled foundation for reward, explanation, and accountability. From its origins in cooperative game theory, the Shapley value has proven remarkably versatile, finding applications in domains its creator never envisioned, from machine learning explainability [24] to sports analytics [22, 23] to climate policy [46]. Its axiomatic foundation gives it normative force that purely ad-hoc attribution methods lack.

Together, these concepts shift our focus from blaming individuals to designing systems [4]. They remind us that better outcomes rarely come from better intentions alone; they come from better rules, better incentives, and better understanding of strategic interaction. When systems produce bad outcomes, the question is not usually "who is at fault?" but rather "what incentives created this equilibrium?" and "how should value be attributed to enable better equilibria?"

In a world increasingly shaped by complex, adaptive systems (markets, platforms, and machine learning models) this shift is not optional. It is essential. As AI systems become more autonomous and consequential, understanding their strategic implications grows more urgent [41]. As organizations become more distributed and collaborative, fair attribution becomes more critical. As global challenges like climate change require unprecedented cooperation [46], mechanism design informed by game theory becomes indispensable [44].

The mathematical elegance of Nash equilibrium and Shapley value should not obscure their profound practical importance. These are not merely theoretical constructs but tools for understanding and reshaping the strategic systems that govern modern life. Their continued relevance that persists more than 70 years after their introduction, testifies to the depth of the problems they address and the power of the insights they provide.

References

- [1] Myerson, R. B. (1991). *Game Theory: Analysis of Conflict*. Harvard University Press.
- [2] Nash, J. (1950). Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences*, 36(1), 48-49.
- [3] Shapley, L. S. (1953). A value for n-person games. In *Contributions to the Theory of Games* (Vol. 2, pp. 307-317). Princeton University Press.

- [4] Maskin, E. (2008). Mechanism design: How to implement social goals. *American Economic Review*, 98(3), 567-576.
- [5] Dwork, C., Hardt, M., Pitassi, T., Reingold, O., & Zemel, R. (2012). Fairness through awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference* (pp. 214-226).
- [6] Barrett, S. (1994). Self-enforcing international environmental agreements. *Oxford Economic Papers*, 46, 878-894.
- [7] Daskalakis, C., Goldberg, P. W., & Papadimitriou, C. H. (2009). The complexity of computing a Nash equilibrium. *SIAM Journal on Computing*, 39(1), 195-259.
- [8] Fudenberg, D., & Levine, D. K. (1998). *The Theory of Learning in Games*. MIT Press.
- [9] Maynard Smith, J., & Price, G. R. (1973). The logic of animal conflict. *Nature*, 246(5427), 15-18.
- [10] Perdomo, J. C., Zrnic, T., Mendler-Dünner, C., & Hardt, M. (2020). Performative prediction. In *International Conference on Machine Learning* (pp. 7599-7609). arXiv:2002.06673
- [11] Jiang, R., Chiappa, S., Lattimore, T., György, A., & Kohli, P. (2019). Degenerate feedback loops in recommender systems. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society* (pp. 383-390). arXiv:1902.10730
- [12] Hardt, M., Megiddo, N., Papadimitriou, C., & Wootters, M. (2016). Strategic classification. In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science* (pp. 111-122). arXiv:1506.06980
- [13] Kerr, S. (1975). On the folly of rewarding A, while hoping for B. *Academy of Management Journal*, 18(4), 769-783.
- [14] Campbell, D. T. (1979). Assessing the impact of planned social change. *Evaluation and Program Planning*, 2(1), 67-90.
- [15] Geirhos, R., Jacobsen, J. H., Michaelis, C., Zemel, R., Brendel, W., Bethge, M., & Wichmann, F. A. (2020). Shortcut learning in deep neural networks. *Nature Machine Intelligence*, 2(11), 665-673. arXiv:2004.07780
- [16] Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., & Fergus, R. (2013). Intriguing properties of neural networks. arXiv:1312.6199
- [17] Vosoughi, S., Roy, D., & Aral, S. (2018). The spread of true and false news online. *Science*, 359(6380), 1146-1151.

- [18] Bail, C. A., Argyle, L. P., Brown, T. W., Bumpus, J. P., Chen, H., Hunzaker, M. F., ... & Volfovsky, A. (2018). Exposure to opposing views on social media can increase political polarization. *Proceedings of the National Academy of Sciences*, 115(37), 9216-9221.
- [19] Dalessandro, B., Perlich, C., Stitelman, O., & Provost, F. (2012). Causally motivated attribution for online advertising. In *Proceedings of the Sixth International Workshop on Data Mining for Online Advertising and Internet Economy* (pp. 1-9).
- [20] Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5), 206-215. arXiv:1811.10154
- [21] Li, H., & Kannan, P. K. (2014). Attributing conversions in a multichannel online marketing environment. *Journal of Marketing Research*, 51(1), 40-56.
- [22] Cervone, D., D'Amour, A., Bornn, L., & Goldsberry, K. (2016). A multiresolution stochastic process model for predicting basketball possession outcomes. *Journal of the American Statistical Association*, 111(514), 585-599. arXiv:1408.0777
- [23] Peña, J. L., & Touchette, H. (2012). A network theory analysis of football strategies. arXiv:1206.6904
- [24] Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems* (pp. 4765-4774). arXiv:1705.07874
- [25] Regulation B (Equal Credit Opportunity Act), 12 CFR Part 1002 (2021). Consumer Financial Protection Bureau.
- [26] Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why should I trust you?": Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 1135-1144). arXiv:1602.04938
- [27] Zech, J. R., Badgeley, M. A., Liu, M., Costa, A. B., Titano, J. J., & Oermann, E. K. (2018). Variable generalization performance of a deep learning model to detect pneumonia in chest radiographs. *JAMA Network Open*, 1(6), e181684. arXiv:1803.10988
- [28] Gururangan, S., Swayamdipta, S., Levy, O., Schwartz, R., Bowman, S. R., & Smith, N. A. (2018). Annotation artifacts in natural language inference data. arXiv:1803.02324
- [29] Goodman, B., & Flaxman, S. (2017). European Union regulations on algorithmic decision-making and a "right to explanation". *AI Magazine*, 38(3), 50-57. arXiv:1606.08813
- [30] Lundberg, S. M., Erion, G. G., & Lee, S. I. (2018). Consistent individualized feature attribution for tree ensembles. arXiv:1802.03888

- [31] Covert, I., Lundberg, S. M., & Lee, S. I. (2020). Understanding global feature contributions with additive importance measures. In *Advances in Neural Information Processing Systems* (Vol. 33, pp. 17212-17223). arXiv:2004.00668
- [32] Frye, C., de Mijolla, D., Begley, T., Cowton, L., Stanley, M., & Feige, I. (2020). Shapley explainability on the data manifold. arXiv:2006.01272
- [33] Sculley, D., Holt, G., Golovin, D., Davydov, E., Phillips, T., Ebner, D., ... & Dennison, D. (2015). Hidden technical debt in machine learning systems. In *Advances in Neural Information Processing Systems* (pp. 2503-2511).
- [34] Maleki, S., Tran-Thanh, L., Hines, G., Rahwan, T., & Rogers, A. (2013). Bounding the estimation error of sampling-based Shapley value approximation. arXiv:1306.4265
- [35] Nash, J. F. (1950). The bargaining problem. *Econometrica*, 18(2), 155-162.
- [36] Shapley, L. S., & Shubik, M. (1954). A method for evaluating the distribution of power in a committee system. *American Political Science Review*, 48(3), 787-792.
- [37] Amodei, D., Olah, C., Steinhardt, J., Christiano, P., Schulman, J., & Mané, D. (2016). Concrete problems in AI safety. arXiv:1606.06565
- [38] Sharif, M., Bhagavatula, S., Bauer, L., & Reiter, M. K. (2016). Accessorize to a crime: Real and stealthy attacks on state-of-the-art face recognition. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security* (pp. 1528-1540).
- [39] Lum, K., & Isaac, W. (2016). To predict and serve? *Significance*, 13(5), 14-19.
- [40] Raghavan, M., Barocas, S., Kleinberg, J., & Levy, K. (2020). Mitigating bias in algorithmic hiring. *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 469-481. arXiv:1906.09208
- [41] Vorobeychik, Y., & Kantarcioglu, M. (2018). Adversarial machine learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 12(3), 1-169. arXiv:1802.06552
- [42] Tramèr, F., Kurakin, A., Papernot, N., Goodfellow, I., Boneh, D., & McDaniel, P. (2017). Ensemble adversarial training: Attacks and defenses. arXiv:1705.07204
- [43] Nisan, N., Roughgarden, T., Tardos, E., & Vazirani, V. V. (2007). *Algorithmic Game Theory*. Cambridge University Press.
- [44] Stavins, R. N. (2008). Addressing climate change with a comprehensive US cap-and-trade system. *Oxford Review of Economic Policy*, 24(2), 298-321.
- [45] Gillespie, T. (2018). *Custodians of the Internet: Platforms, Content Moderation, and the Hidden Decisions That Shape Social Media*. Yale University Press.

[46] Lange, A., Vogt, C., & Ziegler, A. (2007). On the importance of equity in international climate policy. *Energy Economics*, 29(3), 545-559.

[47] Harsanyi, J. C., & Selten, R. (1988). *A General Theory of Equilibrium Selection in Games*. MIT Press.

[48] Mertens, J. F., & Neyman, A. (1981). Stochastic games. *International Journal of Game Theory*, 10(2), 53-66.

[49] Simon, H. A. (1982). *Models of Bounded Rationality*. MIT Press.

[50] Camerer, C. F. (2003). *Behavioral Game Theory*. Princeton University Press. Xiv preprint arXiv:1206.6904*.