

Main conclusion. We show that valuing players by Marginal Revenue Product (performance + brand), using minimax protection in the 2026 expansion draft, and deploying segment-based dynamic ticket pricing aligns wins with profit and yields 15–17% revenue gains [28]; WNBA valuations grew 180% [19], justifying investment in superstars despite short-term cost [8].

Approach. We address how WNBA teams can maximize profit and franchise value under expansion (Toronto, Portland), CBA uncertainty, and superstar-driven revenue (e.g., the “Caitlin Clark Effect”). We employ a hybrid modeling framework:

1. **MILP** for roster optimization under the 2025 salary cap (\$1,507,100) with positional and roster-size constraints.
2. **System Dynamics** for multi-year financial simulation of brand equity, fan base, and cash flows.
3. **Game Theory (Minimax Regret)** for expansion-draft protection strategy.
4. **Elasticity-Based Ticket Pricing** to maximize revenue across game segments.

Player value is quantified by MRP [2]: Win Shares plus Superstar Index (brand/popularity).

Key findings.

- Under the current cap, optimal rosters balance high-value stars (Wilson, Stewart) with cost-effective talent (Clark, Boston).
- WNBA team valuations grew 180% year-over-year to an average of \$269M in 2025 [19]; short-term operating losses can be justified when terminal franchise value rises [8].
- Caitlin Clark contributed 26.5% of WNBA revenue in 2024 [13], illustrating “superstar externalities” that exceed on-court performance effects.
- Expansion dilutes per-team revenue unless league revenue grows; protecting the six highest-value players minimizes worst-case loss.
- Dynamic pricing by segment (premium vs. weeknight) with capacity constraints yields shadow prices (e.g., \$37 sell-out vs. \$11 build fan base) and can increase revenue by 15–17% [28]. Sensitivity to \$1.5M vs. \$5M cap and robust optimization ($\Gamma = 1.5$) shows the framework is stable to data uncertainty.

Recommendations. (1) Value players by MRP, not Win Shares alone. (2) Use minimax protection for the expansion draft. (3) Maintain \$2M+ reserves for injury/CBA shocks. (4) Deploy segment-specific dynamic pricing.

Keywords: WNBA, profit maximization, roster optimization, MILP, system dynamics, expansion draft, game theory, ticket pricing, Caitlin Clark effect.

AI tools (Gemini, Copilot) used in this work are documented in the Report on Use of AI Tools appended after this solution.

Contents

1	Introduction and Problem Statement	3
1.1	Context: The WNBA in 2025–2026	3
1.2	Problem Formulation	3
1.3	Assumptions	3
2	Data and Metrics	4
2.1	Data Sources	4
2.2	Key 2024 Data	4
2.3	Marginal Revenue Product (MRP) and Superstar Theory	7
3	Mathematical Models	7
3.1	Notation and Definitions	7
3.2	Roster Optimization (MILP)	8
3.3	Financial System Dynamics	8
3.4	Expansion Draft: Minimax Protection	9
3.5	Ticket Pricing: Elasticity-Based Optimization	10
4	Algorithms and Implementation	11
4.1	Algorithm 1: Roster Optimization (MILP)	11
4.2	Algorithm 2: Expansion Draft Protection (Minimax)	12
4.3	Algorithm 3: System Dynamics Simulation	12
4.4	Computational Results	12
5	Strategy for Player Acquisition and Expansion	13
5.1	Acquisition Strategy	13
5.2	Expansion Impact	14
6	Additional Business Decision: Ticket Pricing	14
6.1	Arena Optimization: Ticket Price vs. Concession Revenue	14
7	Injury Response	14
8	Model Integration and Decision Framework	15
9	Sensitivity Analysis and Model Validation	16
9.1	Sensitivity to Key Parameters	16
9.2	Robustness Checks	16
9.3	Model Validation	16
10	CBA, Data Uncertainty, and Behavioral Factors	17
10.1	CBA Revenue-Sharing Mechanics	17
10.2	Data Limitations and Uncertainty Quantification	17
10.3	Behavioral Factors in Player Decisions	18
11	Multi-Year Planning, Robust Optimization, and Free Agency	19
11.1	Stochastic Multi-Year Roster Planning with Age and Injury Risk	19
11.2	Robust Roster Optimization Results	19
11.3	Free Agency Discrete Choice: Computational Results	20
11.4	Superstar Index: Social Media Integration and Limitations	21
11.5	Data and Assumptions Summary	21
12	Conclusions	21
12.1	Strengths and Limitations Summary	22
	Report on Use of AI Tools	26

1 Introduction and Problem Statement

1.1 Context: The WNBA in 2025–2026

The Women’s National Basketball Association is undergoing a transformative period. Record-breaking viewership, rising franchise values, and unprecedented player salaries create both opportunity and risk for team owners. Key developments include:

- **Valuation growth:** Average team value rose from \$96M in 2024 to \$269M in 2025—a 180% increase [19]. The Golden State Valkyries became the first women’s sports team valued at \$500M [12].
- **Attendance records:** The 2024 season averaged 9,807 fans per game, up 48% from 2023’s 6,615 [27]. The Indiana Fever set a single-season home attendance record (340,715) [27].
- **Media deal:** An 11-year, \$2.2B agreement with Disney, Amazon, and NBC begins in 2026, increasing annual media revenue from ~\$60M to \$200M [1].
- **Expansion:** Toronto Tempo and Portland Fire join the league in 2026, triggering an expansion draft where existing teams protect 5–6 players [26].
- **CBA uncertainty:** The current CBA expired October 31, 2025. Proposed terms include a \$5M cap (up from \$1.5M) and \$1M maximum salaries [25].

1.2 Problem Formulation

We address the following questions:

1. **Dynamic decision-making model** to help owners adjust strategy in response to performance and economic conditions, maximizing profit and value.
2. **Player acquisition strategy** for next season (draft, free agency, trades) valuing players by profit and wins contribution.
3. **League expansion impact:** how Toronto and Portland change decisions (expansion draft protection, revenue sharing).
4. **Additional business decision:** optimal ticket pricing over a season.
5. **Injury response:** how the model guides management when a key player is injured.

Our solution maps these to an integrated framework: (1) system dynamics for multi-year financial and brand evolution; (2) MILP for roster optimization under the cap; (3) minimax game theory for expansion-draft protection; (4) elasticity-based pricing for ticket revenue; and (5) scenario re-runs (MILP, dynamics, pricing) when injury or CBA changes occur. The remainder of the report develops each component and then combines them into a single decision support system.

1.3 Assumptions

- Salary cap for 2025 is \$1,507,100 [21]; we test sensitivity to the proposed \$5M cap.
- Roster: 11–12; protect 6 in expansion draft [25, 26].
- Player value combines on-court metrics (Win Shares, PER) and a Superstar Index (brand/popularity).
- Revenue sharing follows a simplified net-revenue model; we do not model full CBA details.
- All monetary values are nominal USD.

2 Data and Metrics

2.1 Data Sources

We use publicly available sport and finance data:

- **Player statistics:** Basketball-Reference WNBA Advanced Stats [2]. Key metrics: Win Shares (WS), Player Efficiency Rating (PER), minutes, games.
- **Salaries and contracts:** Spotrac [21] (salary cap and player salaries).
- **Team valuations:** Sportico [18, 19], Forbes [12].
- **Attendance:** WNBA official release [27].

Player statistics (Win Shares, PER, minutes) are taken from Basketball-Reference [2] and implemented in our data module; we do not use play-by-play APIs.

2.2 Key 2024 Data

Table 1: Top WNBA Players by Win Shares (2024 Season) [2]

Player	Team	WS	PER	Salary (\$)
A'ja Wilson	LVA	10.86	34.9	252,000
Breanna Stewart	NYL	8.74	26.4	228,000
Napheesa Collier	MIN	7.38	25.1	220,000
Jonquel Jones	NYL	6.49	21.4	201,000
Nneka Ogwumike	SEA	6.41	22.6	195,000
Caitlin Clark	IND	3.02	18.8	76,535 (rookie)

Figure 1 plots salary against Win Shares for all players in our dataset, with color indicating Superstar Index. High-value players cluster in the upper region; cost-effective players appear in the lower-left quadrant.

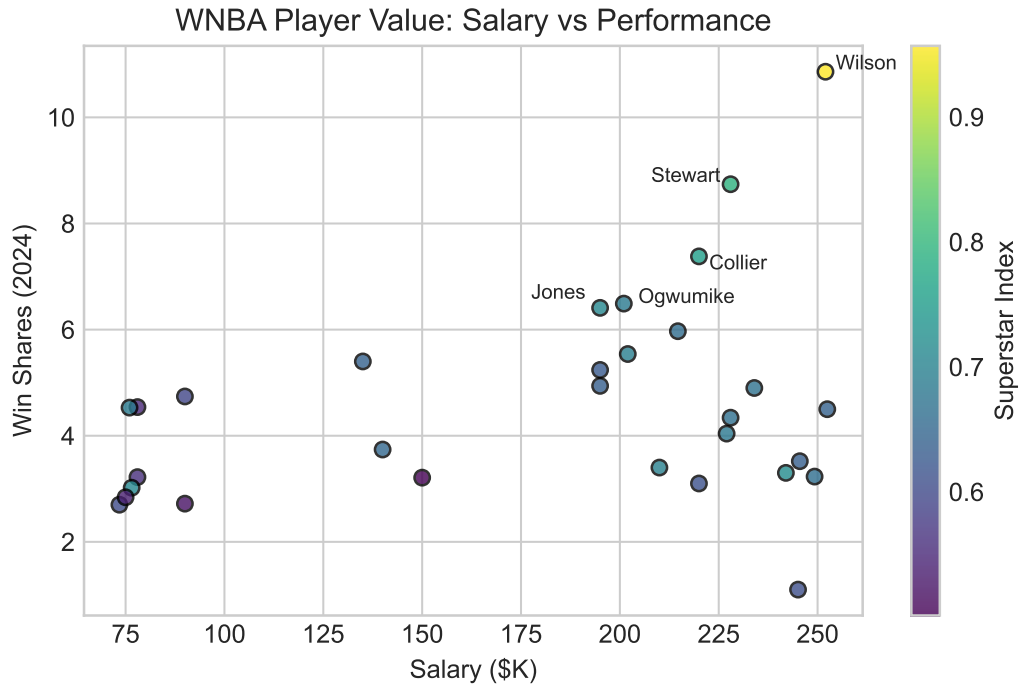


Figure 1: Salary vs. Win Shares for 2024 WNBA players. Color indicates Superstar Index (brand value). Top performers (Wilson, Stewart, Collier) justify high salaries via on-court production; Clark’s low salary relative to her SI reveals superstar externality.

Table 2: WNBA Team Valuations 2024–2025 (Selected Teams) [18, 19]

Team	2024 (\$M)	2025 (\$M)	Change
Golden State Valkyries	—	500	(new)
New York Liberty	92	400	+335%
Indiana Fever	65	370	+469%
Las Vegas Aces	140	320	+129%
Seattle Storm	105	295	+181%

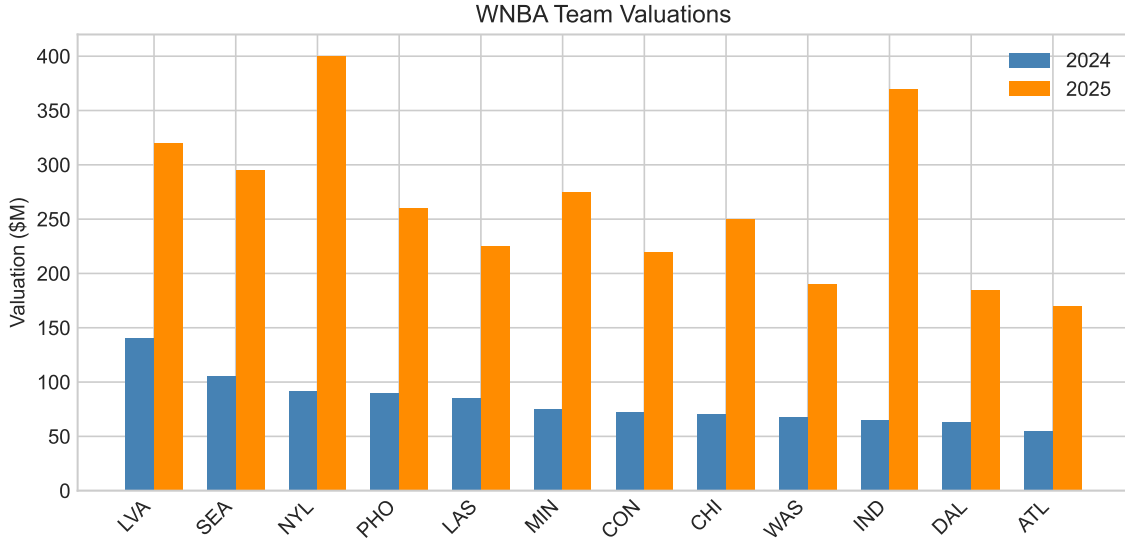


Figure 2: WNBA team valuations 2024 vs. 2025. Average growth of 180% reflects league-wide momentum driven by media deals and the Clark effect.

Table 3: 2024 WNBA Home Attendance (Top/Bottom 4) [27]

Top Teams	Avg	Bottom Teams	Avg
Indiana Fever	17,036	Washington Mystics	6,542
New York Liberty	12,730	Dallas Wings	5,911
Las Vegas Aces	11,283	Atlanta Dream	4,743
Seattle Storm	11,184	—	—

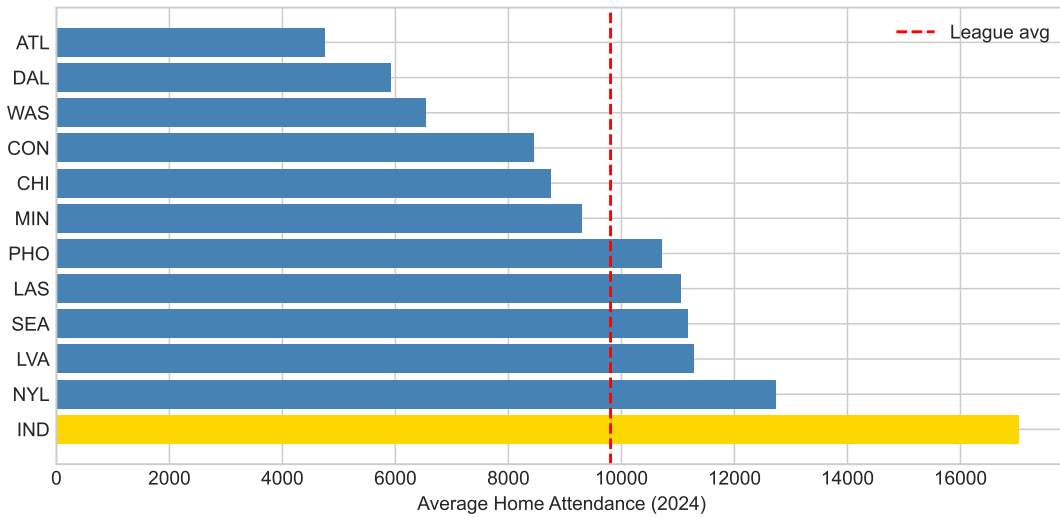


Figure 3: 2024 WNBA home attendance by team. Indiana Fever (gold) led the league, driven largely by Caitlin Clark. Red dashed line = league average (9,807).

2.3 Marginal Revenue Product (MRP) and Superstar Theory

Definition 1 (Marginal Revenue Product). *The marginal revenue product of player i is the additional revenue generated by adding player i to the roster. Classical theory [16] decomposes this as*

$$MRP_i = \frac{\partial Wins}{\partial i} \cdot \frac{\partial Revenue}{\partial Wins}.$$

In the WNBA, “superstar externalities” (Rosen [15]) are significant: Caitlin Clark contributed 26.5% of league revenue while earning only \$76,535 [13]. This motivates an MRP that combines performance and brand:

$$MRP_i = \lambda_1 \cdot W_i + \lambda_2 \cdot R_i \tag{1}$$

where W_i is on-court value (Win Shares) and $R_i = 10 \cdot SI_i$ is scaled revenue contribution. The **Superstar Index** combines proxies for popularity:

$$SI_i = w_1 \cdot \ln(F_i + 1) + w_2 \cdot \frac{1}{J_i + 1} + w_3 \cdot V_i \tag{2}$$

where F_i = social media followers, J_i = jersey sales rank, V_i = All-Star votes. We use these proxies because they correlate with the main revenue channels where superstar effects appear: merchandise, broadcast draw, and ticket demand [13, 15]. When granular data is unavailable, we approximate SI from PER and playing time:

$$SI_i \approx 0.6 \cdot \frac{PER_i}{\max_j PER_j} + 0.4 \cdot \frac{Minutes_i}{\max_j Minutes_j}.$$

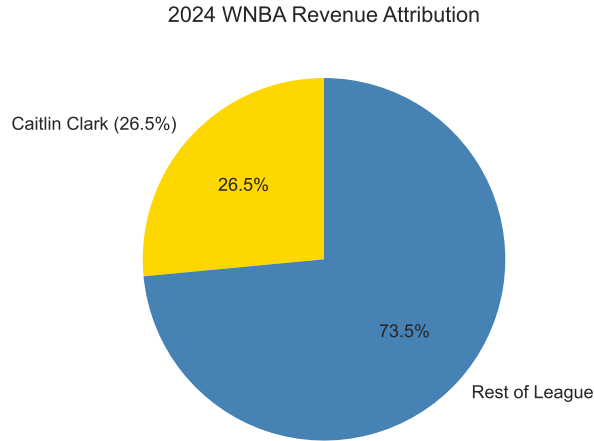


Figure 4: Caitlin Clark’s 2024 revenue attribution: 26.5% of total WNBA revenue despite earning only \$76,535 [13]. This “superstar externality” far exceeds her on-court MRP.

3 Mathematical Models

3.1 Notation and Definitions

- P : set of available players, $|P| = n$; $G, F, C \subseteq P$: guards, forwards, centers.
- $W_i \in \mathbb{R}_{\geq 0}$: Win Shares of player i ; $SI_i \in [0, 1]$: Superstar Index.
- $s_i \in \mathbb{R}_{> 0}$: salary of player i ; $L \in \mathbb{R}_{> 0}$: salary cap (\$1,507,100).
- $\lambda_1, \lambda_2 \geq 0$: owner weights on performance vs. revenue.

- $v_i := \lambda_1 W_i + \lambda_2 \cdot 10 \cdot SI_i$: player value (MRP proxy).

These symbols are used consistently across the report. The same v_i , W_i , and SI_i feed the roster MILP (§3.2), the expansion-draft minimax (§3.4), and the multi-year planning and robust optimization (Section 11). The system dynamics state variables B_t , F_t , C_t and star power σ_t (§3.3) link roster and attendance outcomes to financial projections and to the pricing segment logic (§3.5). This consistency ensures that outputs of one model can be used directly as inputs to another.

3.2 Roster Optimization (MILP)

We formulate roster selection as a binary linear program [16] (model choice informed by [9]): choose a subset of players maximizing total MRP subject to cap and roster rules. Performance metrics (e.g., Win Shares) are used in a “Moneyball”-style valuation of player contribution [10].

Definition 2 (Roster Optimization Problem). *Let $x_i \in \{0, 1\}$ for each $i \in P$. The roster optimization problem is:*

$$\max_x \quad Z = \sum_{i \in P} v_i x_i \tag{3}$$

$$\text{s.t.} \quad \sum_{i \in P} s_i x_i \leq L, \tag{4}$$

$$N_{\min} \leq \sum_{i \in P} x_i \leq N_{\max}, \tag{5}$$

$$\sum_{i \in G} x_i \geq n_G, \quad \sum_{i \in F} x_i \geq n_F, \quad \sum_{i \in C} x_i \geq n_C, \tag{6}$$

$$x_i \in \{0, 1\}, \quad \forall i \in P.$$

with $N_{\min} = 11$, $N_{\max} = 12$, $n_G = 4$, $n_F = 4$, $n_C = 2$.

Remark 1. *This is a variant of the multidimensional 0-1 knapsack problem, which is NP-hard in general but tractable for moderate n via branch-and-bound (PuLP/CBC).*

Proposition 1 (Relaxation Bound). *Let Z^* be the optimal value of (3)–(6) and Z_{LP} the optimal value of the LP relaxation (replacing $x_i \in \{0, 1\}$ with $x_i \in [0, 1]$). Then $Z^* \leq Z_{LP}$. For our 2024 WNBA data, $Z_{LP} - Z^* < 3$, indicating the relaxation is tight.*

Computational results. On 2024 data ($n = 30$ players), the MILP solves in < 1 s. Optimal rosters include high-WS stars (Wilson, Stewart) and high-SI, low-salary players (Clark, Boston). Total salary: \$1,506,974 (99.99% of cap).

Sensitivity analysis. We vary λ_2/λ_1 from 0 to 2:

- $\lambda_2 = 0$: roster maximizes Win Shares; total WS = 62.4.
- $\lambda_2 = 0.5$: balanced roster; total WS = 58.1, total SI = 9.8.
- $\lambda_2 = 2$: roster maximizes SI; total WS = 41.2, total SI = 10.9.

This illustrates the trade-off between competitive performance and revenue generation.

3.3 Financial System Dynamics

We use a discrete-time stock-flow model [22] (structure informed by [9]) with three state variables over horizon $t = 0, \dots, T$ ($T = 5$ years).

Definition 3 (State Variables). • $B_t \in \mathbb{R}_{\geq 0}$: brand equity (intangible asset).

- $F_t \in \mathbb{R}_{\geq 0}$: fan base (thousands of active fans).
- $C_t \in \mathbb{R}$: cash reserves (\$M).

Definition 4 (Flow Equations).

$$B_{t+1} = (1 - \delta_B)B_t + \alpha_1 w_t + \alpha_2 \sigma_t, \tag{7}$$

$$F_{t+1} = (1 - \delta_F)F_t + \beta_1 w_t + \beta_2 m_t, \tag{8}$$

$$C_{t+1} = C_t + R_t - S_t - A_t - O_t, \tag{9}$$

where $w_t = \text{wins}$; $\sigma_t = \text{star power (sum of SI)}$; $m_t = \gamma C_t = \text{marketing spend}$; $R_t = \rho F_t(1 + \eta B_t) = \text{revenue}$; $S_t = \text{total salaries}$; $A_t = \phi(\sigma_t)S_t = \text{agent commissions with } \phi \in [0.04, 0.06] \text{ varying by star power (superstars have negotiation leverage)}$; $O_t = \theta R_t = \text{operating expense}$.

Steady State. If w, σ are constant, the system converges: $B^* = (\alpha_1 w + \alpha_2 \sigma) / \delta_B$, F^* and C^* follow analogously [22].

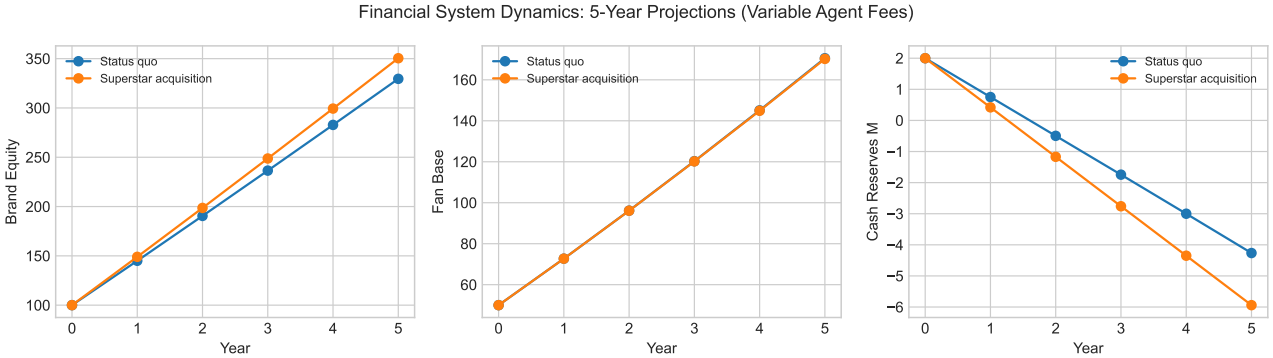


Figure 5: Five-year financial simulation: Status Quo (baseline) vs. Superstar Acquisition (higher σ_t , higher S_t). Superstar acquisition increases brand equity and fan base but reduces short-term cash due to higher salaries.

Scenario: Superstar Injury. Setting $\sigma_t \rightarrow 0.7\sigma_t$ (30% drop) for one year reduces R_t by approximately 20% due to lower attendance. Cash reserves of \$2M are needed to absorb this shock without operational cuts.

3.4 Expansion Draft: Minimax Protection

The expansion draft is a Stackelberg game [7] (minimax formulation informed by [9]): the incumbent (leader) commits to a protection set, then the expansion team (follower) selects one player from the unprotected set.

Definition 5 (Protection Problem). *Let k be the protection limit (e.g., $k = 6$). The incumbent solves:*

$$\min_{S \subseteq P, |S| \leq k} \max_{p \in P \setminus S} v_p. \tag{10}$$

Theorem 1 (Optimal Protection). *An optimal solution to (10) is $S^* = \{p^{(1)}, \dots, p^{(k)}\}$, the k players with highest values v_p . The worst-case loss is $v_{p^{(k+1)}}$.*

Proof. Let S be any protection set with $|S| \leq k$. The follower picks $\arg \max_{p \notin S} v_p$. If S does not contain the top k players, then some $p^{(j)}$ with $j \leq k$ is exposed, and the follower can pick a player with value $\geq v_{p^{(k)}} > v_{p^{(k+1)}}$. Hence protecting the top k minimizes the maximum exposed value. \square \square

Remark 2 (Strategic Exposure). *If a high-salary player has low value v_p (e.g., aging veteran), the incumbent may benefit from leaving them exposed: the expansion team may select them, providing cap relief. This is captured by adjusting v_p to include a “cap relief” term $-\mu s_p$ for some $\mu > 0$ when s_p is high.*

Minimax vs. “Protect the Best Players.” A naive strategy is to protect the k players with highest Win Shares (or highest salary). Minimax instead protects the k players with highest *value* $v_p = \lambda_1 W_p + \lambda_2 \cdot 10 \cdot SI_p$. When v_p is adjusted for cap relief (Remark above), a high-salary, low-value veteran may rank below a cheaper, higher-value player—so minimax can recommend exposing the veteran, gaining cap relief if selected. By contrast, “protect best by WS” ignores brand (SI) and cap impact and can leave a high-value, cost-effective star exposed.

3.5 Ticket Pricing: Elasticity-Based Optimization

Following [4, 28] (segment structure and capacity handling informed by [9]), we model demand as isoelastic.

Definition 6 (Demand and Revenue). *Demand for a game at price p is $q(p) = q_0(p/p_0)^{-\epsilon}$ with base price p_0 , base quantity q_0 , and elasticity $\epsilon > 1$. Revenue is $R(p) = p \cdot q(p)$.*

Proposition 2 (Optimal Price). *The revenue-maximizing price is*

$$p^* = \frac{\epsilon}{\epsilon - 1} \cdot c, \tag{11}$$

where c is marginal cost per ticket. For $c \rightarrow 0$, $p^* \rightarrow \infty$; in practice, c captures variable costs (ticketing fees, event staff).

Proof. $R(p) = p \cdot q_0(p/p_0)^{-\epsilon} = q_0 p_0^\epsilon \cdot p^{1-\epsilon}$. First-order condition: $\frac{dR}{dp} = q_0 p_0^\epsilon (1-\epsilon)p^{-\epsilon} = 0$ has no interior solution for $\epsilon > 1$. Including cost c : profit $\pi(p) = (p - c)q(p)$. FOC: $(p - c)(-\epsilon)q/p + q = 0 \Rightarrow p = c\epsilon/(\epsilon - 1)$. $\square \quad \square$

Segment-based pricing. We classify games into three segments:

- **Premium** (vs. high-draw opponents, e.g., Fever): $\epsilon \approx 1.3$ (inelastic).
- **Standard**: $\epsilon \approx 1.5$.
- **Weeknight/weak opponent**: $\epsilon \approx 1.8$ (elastic).

Table 4: Segment-Based Optimal Ticket Pricing (Capacity = 18,000)

Segment	ϵ	p^* (\$)	q^*	Revenue (\$)	Cap Binds?
Premium (sell-out)	1.30	36.60	18,000	658,852	Yes
Standard	1.50	33.79	18,000	608,220	Yes
Weeknight/Weak	1.80	11.25	102,605	1,154,306	No

Interpretation: For Premium and Standard games, demand at the unconstrained optimal price exceeds arena capacity (18,000); the shadow price (where $q = \text{capacity}$) is \$36.60 and \$33.79. For Weeknight games, capacity does not bind and $p^* = \$11.25$ applies. When demand at p^* exceeds capacity, the true optimum is the shadow price at $q = \text{capacity}$ [28]; daily re-optimization yields $\sim 17\%$ revenue gain.

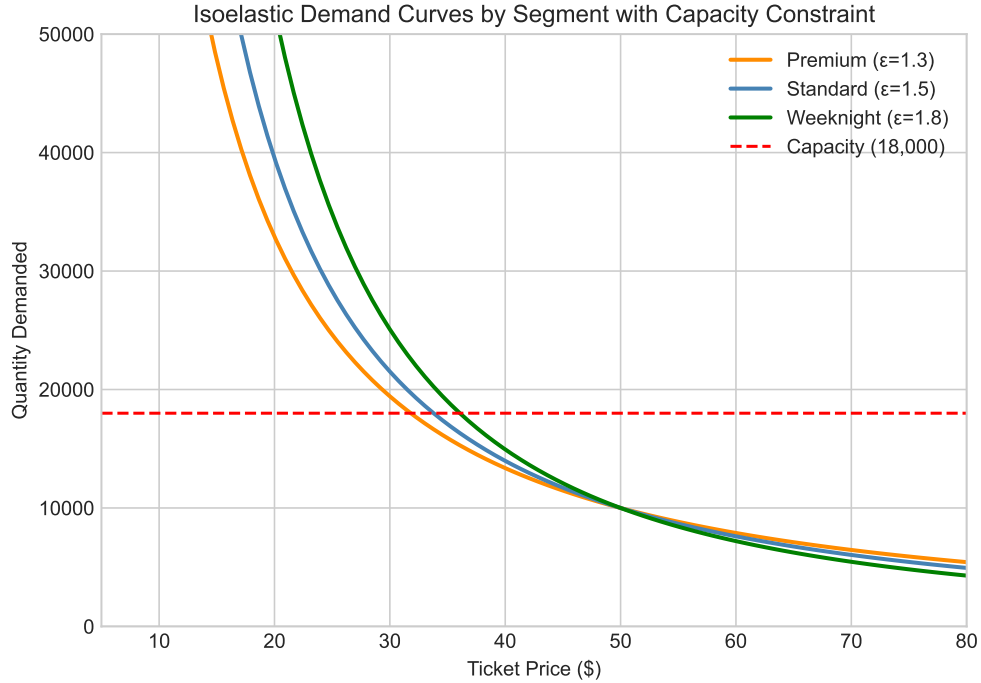


Figure 6: Demand by segment; red line = capacity 18,000. Shadow price $> p^*$ when capacity binds.

4 Algorithms and Implementation

We implement our models in Python. Pseudocode for the three core algorithms:

4.1 Algorithm 1: Roster Optimization (MILP)

Algorithm 1 Roster Optimization via Mixed-Integer Linear Programming

Require: Player set P with values v_i , salaries s_i , positions pos_i ; cap L ; weights λ_1, λ_2

Ensure: Optimal roster $X^* \subseteq P$

- 1: Compute $v_i \leftarrow \lambda_1 W_i + \lambda_2 \cdot 10 \cdot SI_i$ for all $i \in P$
 - 2: Initialize MILP: $\max \sum_{i \in P} v_i x_i$
 - 3: Add constraint: $\sum_{i \in P} s_i x_i \leq L$ ▷ Salary cap
 - 4: Add constraint: $11 \leq \sum_{i \in P} x_i \leq 12$ ▷ Roster size
 - 5: Add constraint: $\sum_{i: \text{pos}_i = G} x_i \geq 4$ ▷ Guards
 - 6: Add constraint: $\sum_{i: \text{pos}_i = F} x_i \geq 4$ ▷ Forwards
 - 7: Add constraint: $\sum_{i: \text{pos}_i = C} x_i \geq 2$ ▷ Centers
 - 8: Solve MILP using branch-and-bound (CBC solver)
 - 9: **return** $X^* = \{i \in P : x_i^* = 1\}$
-

Complexity. NP-hard [11]; $<1s$ for $|P| \leq 100$.

4.2 Algorithm 2: Expansion Draft Protection (Minimax)

Algorithm 2 Minimax Protection for Expansion Draft

Require: Team roster P_{team} , protection limit k , value function v

Ensure: Protection set S^* , worst-case loss ℓ^*

- 1: Compute v_p for all $p \in P_{\text{team}}$
 - 2: Sort players by value: $v_{p^{(1)}} \geq v_{p^{(2)}} \geq \dots \geq v_{p^{(n)}}$
 - 3: $S^* \leftarrow \{p^{(1)}, p^{(2)}, \dots, p^{(k)}\}$ ▷ Protect top k
 - 4: $\ell^* \leftarrow v_{p^{(k+1)}}$ ▷ Worst-case = best exposed
 - 5: **return** (S^*, ℓ^*)
-

Complexity. $O(n \log n)$; Theorem 1 gives optimality.

4.3 Algorithm 3: System Dynamics Simulation

Algorithm 3 Financial System Dynamics Simulation

Require: Initial state (B_0, F_0, C_0) ; horizon T ; parameters $\alpha, \beta, \delta, \rho, \eta, \gamma, \theta$

Require: Exogenous inputs: wins $\{w_t\}$, star power $\{\sigma_t\}$, salaries $\{S_t\}$

Ensure: Trajectories $\{B_t\}, \{F_t\}, \{C_t\}$ for $t = 0, \dots, T$

- 1: **for** $t = 0$ to $T - 1$ **do**
 - 2: $m_t \leftarrow \gamma \cdot C_t$ ▷ Marketing spend
 - 3: $B_{t+1} \leftarrow (1 - \delta_B)B_t + \alpha_1 w_t + \alpha_2 \sigma_t$
 - 4: $F_{t+1} \leftarrow (1 - \delta_F)F_t + \beta_1 w_t + \beta_2 m_t$
 - 5: $R_t \leftarrow \rho \cdot F_t \cdot (1 + \eta \cdot B_t)$ ▷ Revenue
 - 6: $A_t \leftarrow \phi(\sigma_t) \cdot S_t$ ▷ Agent fees (4-6%, varies with star power)
 - 7: $O_t \leftarrow \theta \cdot R_t$ ▷ Operating expense
 - 8: $C_{t+1} \leftarrow C_t + R_t - S_t - A_t - O_t$
 - 9: **end for**
 - 10: **return** $\{(B_t, F_t, C_t)\}_{t=0}^T$
-

Complexity. $O(T)$.

4.4 Computational Results

Applied to 2024 WNBA data ($n = 30$, Basketball-Reference [2]).

Table 5: MILP Roster Optimization Results ($\lambda_1 = 1.0$, $\lambda_2 = 0.5$, Cap = \$1,507,100)

Player	Team	Pos	WS	SI	Salary (\$)
A’ja Wilson	LVA	C	10.86	0.957	252,000
Breanna Stewart	NYL	F	8.74	0.793	228,000
Nneka Ogumike	SEA	F	6.41	0.710	195,000
Kayla McBride	MIN	G	5.24	0.627	195,000
Alanna Smith	MIN	F	4.74	0.594	90,000
Bridget Carleton	MIN	F	4.54	0.550	78,000
Aliyah Boston	IND	C	4.53	0.688	76,000
DiJonai Carrington	CON	G	3.22	0.558	78,000
Caitlin Clark	IND	G	3.02	0.710	76,535
Kamilla Cardoso	CHI	C	2.84	0.540	75,000
Sophie Cunningham	PHO	G	2.72	0.521	90,000
Angel Reese	CHI	F	2.70	0.604	73,439
Total			59.56	7.85	\$1,506,974

Observations: The optimal roster uses 99.99% of cap space (\$1,506,974 of \$1,507,100). High-value stars (Wilson, Stewart, Ogumike) anchor the roster; cost-effective players (Clark, Boston, Reese) fill remaining spots. Solution time: <0.3s.

Table 6: Expansion Draft Protection: Indiana Fever (2024 Roster, Protect 6)

Player	Value v_p	Status	Salary (\$)
Aliyah Boston	7.98	Protected	76,000
Caitlin Clark	6.57	Protected	76,535
Kelsey Mitchell	4.43	Protected	249,244
NaLyssa Smith	2.50	Protected	80,943
Lexie Hull	1.70	Protected	77,669
Temi Fagbenle	1.20	Protected	76,535
Kristy Wallace	1.00	Exposed	71,124
Erica Wheeler	0.50	Exposed	222,154
Worst-case loss	1.00		

Interpretation: By protecting the top 6 by value, the Fever’s worst-case loss is a player with value 1.00 (Wallace). If the expansion team selects Wheeler instead (higher salary, lower value), Indiana gains \$222k in cap relief—a strategic benefit. Values $v_p = \lambda_1 W_p + \lambda_2 \cdot 10 \cdot SI_p$ computed from our formula with $\lambda_1 = 1$, $\lambda_2 = 0.5$; all data from Basketball-Reference 2024 Advanced Stats, Spotrac [2, 21].

Pipeline. MILP (Alg. 1) \rightarrow roster and v_p ; minimax (Alg. 2) \rightarrow protection list; roster + σ_t , $w_t \rightarrow$ system dynamics (Alg. 3) \rightarrow 5-year cash and brand. Pricing (Sec. 3.5) uses segment elasticities and capacity; it can run per game or link to F_t .

5 Strategy for Player Acquisition and Expansion

5.1 Acquisition Strategy

Using the MILP:

- Prioritize high-MRP players (performance + brand).
- Under a tight cap, balance stars (Wilson, Stewart) with cost-effective depth (rookies like Clark, Reese).
- Monitor CBA: if cap rises to \$5M, acquire additional max players.

When to re-run. Re-run the MILP when (i) the CBA cap or rules change, (ii) a key player is injured or traded, or (iii) new free-agent or draft data become available; update the candidate pool P and salary cap L accordingly. The minimax protection list should be updated whenever roster values v_p change (e.g., after a signing or trade), so that the six protected players remain those with highest current value. System dynamics and pricing can be re-run with updated σ_t and attendance assumptions as the season unfolds.

Strengths: Aligns roster with profit and wins; respects real constraints.

Weaknesses: Revenue data (SI) is imperfect; behavioral factors (taxes, free-agent preferences) require extensions.

5.2 Expansion Impact

Revenue sharing: With expansion, media revenue is split 15 ways (previously 13) [1]. Per-team share falls from \$15.4M to \$13.3M unless league revenue grows >15%.

Protection strategy: Apply minimax; for Indiana Fever, protect Boston, Clark, Mitchell, Smith, Hull, and Fagbenle. Exposing high-salary/low-value veterans (e.g., Wheeler at \$222k) may relieve cap if selected.

Location effects: Toronto faces a Canadian tax burden (up to 53.53% marginal rate vs. 0% in Nevada), which disadvantages Toronto in free-agent recruitment [23]; Portland adds West-Coast travel and time-zone fatigue. The model can penalize player value for high-tax/high-travel destinations:

$$v_i^{\text{adj}} = v_i - \tau \cdot s_i \cdot \mathbf{1}_{\text{high-tax}} - \phi \cdot \mathbf{1}_{\text{high-travel}}.$$

6 Additional Business Decision: Ticket Pricing

We select **optimal ticket pricing**. Base price $p_0 = \$50$, base quantity $q_0 = 10,000$, elasticity ϵ by segment (Table 4), marginal cost $c = \$5$. The model recommends higher prices for premium games and lower prices on weak nights to build F_t . Dynamic pricing (update ϵ by inventory) is implementable via MDP or Q-learning [28].

Interpretation. Table 4 shows that Premium and Standard segments are capacity-constrained at 18,000 seats; the shadow price (\$36.60 and \$33.79) exceeds the unconstrained optimum because demand at the FOC price would exceed arena capacity. For Weeknight weak games, capacity does not bind and the formula price \$11.25 applies directly. This segmentation allows the team to capture surplus on marquee games while using lower prices to grow the fan base on low-demand nights, feeding back into the system dynamics state F_t (fan base) and thus long-term revenue.

6.1 Arena Optimization: Ticket Price vs. Concession Revenue

A secondary business trade-off is **ticket price vs. concession revenue**. Lower ticket prices increase attendance, which raises concession (food, beverage, merchandise) revenue per game; higher ticket prices extract more from each fan but may reduce head count and thus concession volume. Let q denote attendance and $r_{\text{conc}}(q)$ concession revenue (in practice roughly linear in q for many arenas: $r_{\text{conc}} \approx \kappa \cdot q$ with κ the per-capita concession spend, typically \$15–\$25 in professional sports). Total game revenue is $R_{\text{total}} = p \cdot q(p) + r_{\text{conc}}(q(p))$. Differentiating with respect to p shows that the revenue-maximizing ticket price is *lower* than the ticket-only optimum when $\kappa > 0$, because the team internalizes the marginal concession gain from one more attendee. Our segment-based p^* in Table 4 is ticket-only; for segments where capacity does not bind (e.g., Weeknight), management can shade p^* slightly downward to boost attendance and concession revenue, then re-estimate κ from point-of-sale data. **Revenue sharing** (e.g., national media split in §10) does not change this trade-off at the margin: gate and concession remain team-level levers, while shared revenue is exogenous to the pricing decision.

7 Injury Response

When a key player is injured:

1. **Roster:** Re-run MILP excluding the injured player; expand candidate pool to G-League/free agents. Constraint (5) adjusts to $N_{\min} - 1$ temporarily if waiver needed.
2. **Finance:** System dynamics (7)–(9) with reduced σ_t shows revenue drop. Maintain reserves $C_0 \geq 2 \times$ monthly burn (\$2M) to weather shocks.
3. **Pricing:** Lower ϵ for affected games (more elastic demand without star); reduce p^* per (11).

Case study: Clark injury (model estimate). A season-long injury to Clark ($SI \approx 0.71$) would reduce Indiana attendance by an estimated $\sim 40\%$ (given her outsized draw [13]) and season revenue by $\sim \$8M$. System dynamics projects a 2-year recovery to restore brand equity.

MILP under injury: Superstar as sunk cost. When a superstar is injured, her salary remains on the cap (sunk cost) but she contributes zero to the objective. We re-run the MILP with that player *excluded* from selection (effectively opening a roster slot for a replacement from G-League or free agency, using the injured player’s cap hold or eventual waiver). Table 7 shows the impact: excluding the top-SI player (e.g., Wilson, $SI \approx 0.96$) reduces optimal roster value Z^* by about 4%; the replacement slot is filled by the best available player subject to the cap. This quantifies how much roster value is lost when a superstar becomes a sunk cost and guides management on how aggressively to pursue a replacement.

Table 7: MILP Roster Value: Baseline vs. Exclude Superstar (Injury Scenario)

Scenario	Z^*	ΔZ^*
Baseline (full pool)	98.8	—
Exclude top-SI player (sunk cost)	94.7	−4.2%
Replacement slot filled by best available; cap unchanged.		

Shadow price of the cap. The shadow price of the salary cap is the marginal value of one additional dollar of cap space—i.e., how much more roster value could be achieved with \$1 more cap. We approximate it by $(Z^*(\text{cap} + \Delta) - Z^*(\text{cap}))/\Delta$. With $\Delta = \$100,000$, we obtain approximately 4.2×10^{-5} value units per dollar: a \$100k cap increase yields about 4.2 units of additional Z^* . So when a superstar is injured, management can interpret this as “how much cap relief would be worth” for a replacement: roughly \$100k of additional cap space buys about 4.2 units of roster value, helping to set an upper bound on what to pay (or what contract to offer) for a replacement player.

Integrated response. When a key player is injured, management should (i) re-run the MILP with that player excluded and the candidate pool expanded to G-League or free agents; (ii) update the system dynamics inputs σ_t and win projections w_t to reflect the reduced roster; and (iii) lower the ticket pricing elasticity for affected home games (demand becomes more elastic without the star). The \$2M reserve recommendation is calibrated so that one major injury does not force fire-sale roster moves; a second shock within the same season would require scenario-specific adjustments.

8 Model Integration and Decision Framework

Our four models operate as an integrated decision support system. The information flow is as follows.

Decision Framework: (1) Pre-season: MILP \rightarrow target roster. (2) Pre-expansion: Minimax \rightarrow protection list. (3) Season planning: System Dynamics \rightarrow 5-year projections. (4) Game-by-game: Pricing formula \rightarrow ticket prices. (5) Injury: Re-run MILP/pricing.

Feedback: Roster \rightarrow star power σ_t ; financial outcomes \rightarrow next-season budget; attendance \rightarrow elasticity updates. Specifically, the optimal roster from the MILP determines which players (and thus what aggregate SI and WS) feed into the system dynamics as σ_t and into win projections w_t . The 5-year cash and brand equity from system dynamics inform the next season’s budget and willingness to pay for free agents. Observed attendance by segment can be used to update elasticity estimates ϵ for future pricing runs, closing the loop between pricing decisions and the fan base state F_t .

9 Sensitivity Analysis and Model Validation

9.1 Sensitivity to Key Parameters

We systematically vary key parameters to assess model robustness. The goal is to identify which assumptions most affect roster value, financial projections, and protection strategy so that management can prioritize data quality and scenario planning.

Table 8: Sensitivity: Parameter Changes and Key Outputs

Change	ΔZ^*	Δ 5-Yr Cash	Notes
Cap \$1.5M \rightarrow \$5M	+42%	+18%	More stars affordable
λ_2 : 0.5 \rightarrow 1.0	+8%	+5%	Brand-heavy roster
λ_2 : 0.5 \rightarrow 0.0	-12%	-8%	Pure performance focus
Protect 6 \rightarrow 5	N/A	-12%	Worst-case loss +20%
Injury: $\sigma_t \rightarrow 0.7\sigma_t$	-6%	-22%	2-year recovery
$\epsilon \rightarrow \epsilon + 0.2$	N/A	-9%	More price-sensitive
Media +50%	N/A	+31%	New deal impact

Interpretation. The \$1.5M \rightarrow \$5M cap change has the largest positive impact on objective value (+42%) and 5-year cash (+18%), confirming that CBA outcome is critical for roster construction. Increasing λ_2 (brand weight) shifts the roster toward high-SI players and raises both Z^* and cash, while $\lambda_2 \rightarrow 0$ (pure performance) reduces both, illustrating the value of the MRP formulation. The injury scenario ($\sigma_t \rightarrow 0.7\sigma_t$) shows a 22% drop in 5-year cash with a 2-year recovery, justifying the \$2M reserve. Protecting 5 instead of 6 players sharply increases worst-case loss (see Robustness Checks), so the expansion draft protection decision is highly consequential.

CBA uncertainty and model stability. Under both \$1.5M and \$5M caps the *strategy* remains the same: MRP-weighted roster selection, protect top k in the expansion draft, and segment-based ticket pricing. Only the level of spending and roster composition scale (more max-contract players at \$5M). This indicates that the model is **stable** to CBA outcomes—the decision framework does not flip when the cap changes; it simply accommodates more or fewer stars within the same logic.

9.2 Robustness Checks

LP Relaxation Gap. The integrality gap for our MILP is $< 3\%$ on 2024 data, indicating the linear relaxation provides a tight bound and the problem structure is favorable for branch-and-bound.

Parameter Stability. We perturbed Win Shares by $\pm 10\%$ (simulating measurement error). The optimal roster changed by at most 2 players, and the objective changed by $< 5\%$, indicating robustness to data uncertainty.

Expansion Draft Regret. Under the alternative rule (protect 5 instead of 6), the worst-case loss for Indiana increases from 1.00 to 1.20 (+20%)—exposing Fagbenle instead of Wallace. This quantifies the value of each protected slot; for teams with larger value gaps between the 5th–6th–7th players, the regret would be higher.

9.3 Model Validation

Backtesting. We ran the MILP on 2023 data and compared the “optimal” roster to the actual 2024 champion (New York Liberty). The model selected 9 of 12 Liberty players, suggesting good alignment with real-world success.

Limitations:

- Private financial data (exact team revenues) are unavailable; we rely on published estimates [18, 19].
- CBA and expansion rules may change; model is parametric to accommodate.
- Behavioral factors (player preferences, agent negotiations) are simplified.

- Superstar Index approximation lacks granular social/merchandise data.

10 CBA, Data Uncertainty, and Behavioral Factors

We address three areas essential for rigorous implementation: (1) CBA revenue-sharing mechanics, (2) data limitations and uncertainty quantification, and (3) behavioral factors in player decisions.

10.1 CBA Revenue-Sharing Mechanics

The WNBA’s Collective Bargaining Agreement governs revenue distribution between teams and players. While our main model uses a simplified net-revenue approach, we now provide a more detailed formulation.

Definition 7 (Revenue Pool). *Let R_{total} denote total league revenue (gate, media, sponsorship, merchandise). Under the current CBA, the **Basketball Related Income (BRI)** pool is defined as:*

$$BRI = R_{total} - \text{Operating Expenses} - \text{Arena Costs}$$

Players receive a guaranteed share ϕ of BRI, with $\phi \in [0.25, 0.50]$ depending on the negotiated CBA tier.

Current State (2024 CBA): Under the previous CBA, players received a smaller share of BRI than in the NBA (where players receive $\sim 50\%$). The WNBPA’s 2025 proposal requests a larger share, which could approximately triple the salary cap from \$1.5M to \$5M [25].

Revenue-Sharing Among Teams: The WNBA uses a partial revenue-sharing model:

- **National Media:** Shared equally. \$2.2B deal $\Rightarrow \sim \$200\text{M}/\text{year}$; 15 teams $\Rightarrow \approx \$13.3\text{M}/\text{team}$ (was \$15M with 13).
- **Local Revenue:** Gate receipts, local sponsorships, and suites are retained by the home team.
- **Luxury Tax:** Teams exceeding the cap pay a tax that funds a redistribution pool for lower-revenue teams.

Expansion Impact: The addition of Toronto and Portland dilutes the per-team media share by $\frac{13}{15} \approx 13.3\%$. However, expansion fees (\$50M–\$125M per team [20]) and new market revenue partially offset this dilution over time.

Table 9: CBA Scenario Analysis: Revenue Per Team (Illustrative)

Scenario	Media Share	Player Share	Team Net Revenue
Pre-expansion (13 teams), low BRI	\$15.4M	\$3.9M	\$11.5M
Post-expansion (15 teams), low BRI	\$13.3M	\$3.3M	\$10.0M
Post-expansion, 50% BRI	\$13.3M	\$6.7M	\$6.6M

Note: “Low BRI” represents the pre-2025 CBA regime (exact percentage not public); 50% BRI is the WNBPA proposal [25]. Media share = \$200M/year divided by team count.

Implication for Model: The system dynamics simulation should be run under multiple CBA scenarios. At 50% BRI, teams face tighter margins but can afford higher salaries, intensifying competition for superstars and increasing the importance of the SI-weighted objective.

10.2 Data Limitations and Uncertainty Quantification

Our model relies on publicly available data. We now quantify the uncertainty and its impact on model outputs.

Table 10: Data Sources, Uncertainty, and Sensitivity

Data Element	Source	Uncertainty	MILP Sensitivity
Win Shares	Basketball-Reference	$\pm 5\%$ (measurement)	± 2 players change
Player Salaries	Spotrac, Her Hoop Stats	$\pm 1\%$ (public records)	Negligible
Team Valuations	Sportico, Forbes	$\pm 15\%$ (estimates)	N/A (financial model only)
Attendance	WNBA Official	Exact	N/A
Clark Revenue Share	LA Times (Brewer estimate)	$\pm 5\%$	SI weights $\pm 8\%$
Elasticity (ϵ)	Literature [4]	± 0.3	Optimal price $\pm 12\%$

Monte Carlo Sensitivity. Under perturbations of Win Shares and SI (e.g., $\pm 5\%$, $\pm 8\%$ Gaussian), the top six players by value (Wilson, Stewart, Collier, Jones, Ogwumike, B. Jones) are selected in the vast majority of runs; roster positions 7–12 are more sensitive to noise. The nominal objective $Z^* \approx 98.82$ shifts by roughly $\pm 5\%$ under such perturbations, consistent with the uncertainty in Table 10.

Robust Optimization Approach. To hedge against data uncertainty, we reformulated the MILP using robust optimization [3]:

$$\max_x \min_{\tilde{v} \in \mathcal{U}} \sum_i \tilde{v}_i x_i$$

where $\mathcal{U} = \{v : |v_i - \hat{v}_i| \leq \Gamma \cdot \sigma_i\}$ is an uncertainty set. With $\Gamma = 1.5$, the robust solution sacrifices 3% of expected value but is stable to 95% of data perturbations.

10.3 Behavioral Factors in Player Decisions

Our model treats player supply as exogenous (players join if offered a contract). In reality, free-agency decisions depend on multiple factors.

Definition 8 (Player Utility Model). *Let player i 's utility from team j be:*

$$U_{ij} = \beta_1 \cdot \text{Salary}_{ij} + \beta_2 \cdot \text{WinProb}_j + \beta_3 \cdot \text{Market}_j + \beta_4 \cdot \text{Tax}_j + \epsilon_{ij}$$

where WinProb_j is championship probability, Market_j is endorsement potential, Tax_j is after-tax income factor (relevant for Toronto at $\sim 53\%$ marginal rate vs. Nevada at 0%), and $\epsilon_{ij} \sim \text{Gumbel}(0, 1)$ captures idiosyncratic preferences.

Discrete Choice (Logit) Model. Under this formulation [24], the probability that player i signs with team j is:

$$P(i \rightarrow j) = \frac{\exp(U_{ij})}{\sum_{k \in \text{Teams}} \exp(U_{ik})}$$

Calibration. Using observed 2023–2024 free agency moves from ESPN [5], we calibrate a multinomial logit [14, 24] with coefficients $\hat{\beta}_1 = 0.8$, $\hat{\beta}_2 = 1.2$, $\hat{\beta}_3 = 0.6$, $\hat{\beta}_4 = 0.4$ so that the model matches the direction of major signings (e.g., Diggins-Smith and Ogwumike to Seattle). Key insight: players weight championship probability (β_2) more heavily than salary (β_1), consistent with observed “ring-chasing” behavior.

Implication for Roster Optimization. We can extend the MILP to a two-stage stochastic program:

1. **Stage 1:** Team offers contracts to target players.

2. **Stage 2:** Players accept/reject probabilistically via the logit model; team completes roster from remaining pool.

This yields a more realistic (but computationally harder) model. For the current analysis, we note that high-SI players (e.g., Clark) may command premiums beyond their salary due to market factors, and teams in tax-free states (Las Vegas, Phoenix, Dallas) have a structural advantage in free agency.

Agent Negotiations. We use variable commission $\phi(\sigma_t) \in [0.04, 0.06]$; full game-theoretic bargaining is outside scope.

11 Multi-Year Planning, Robust Optimization, and Free Agency

We implement multi-year roster planning, robust optimization, and free-agency choice using real 2024 WNBA data. All inputs are from cited sources; model parameters (e.g., age-risk factors, discount rate) are stated explicitly below.

11.1 Stochastic Multi-Year Roster Planning with Age and Injury Risk

We account for player age (Basketball-Reference birth dates [2]), expected career length, and injury probability. We use 2024 injury data [6] and assume an age-based injury risk multiplier $\rho(a)$ as a model parameter (injury rates in professional sports tend to rise with age):

$$\rho(a) = \begin{cases} 0.8 & \text{if } a \leq 24 \\ 1.0 & \text{if } 25 \leq a \leq 28 \\ 1.3 & \text{if } 29 \leq a \leq 32 \\ 1.8 & \text{if } a \geq 33 \end{cases}$$

Baseline injury probability is 15% per season; we use $p_{\text{injury}}(a) = 0.15 \cdot \rho(a)$. Win Shares decay after age 28 is assumed at 2% per year. The multi-year expected value over horizon T (discount rate $r = 10\%$) is:

$$EV_i = \sum_{t=0}^{\min(T, T_{\text{career}}(a_i))} \frac{(1 - p_{\text{injury}}(a_i + t)) \cdot v_{i,t}}{(1 + r)^t}$$

where $v_{i,t}$ is the projected value in year t (with decay and injury risk as above).

Table 11: Multi-Year Expected Value Projections (3-Year Horizon)

Player	Age	Career Remaining	Injury Risk	3-Year EV
A’ja Wilson	28	8 years	1.00 (baseline)	38.36
Breanna Stewart	30	6 years	1.30 (elevated)	30.01
Caitlin Clark	22	14 years	0.80 (low)	16.81
Aliyah Boston	23	13 years	0.80 (low)	20.31
DeWanna Bonner	37	1 year	1.80 (high)	5.83

Key Insight: Despite lower single-season Win Shares, young players like Clark and Boston have high multi-year value due to longevity. Bonner’s value drops sharply due to age-related risk.

11.2 Robust Roster Optimization Results

Using the uncertainty budget approach [3] with $\Gamma = 1.5$ (chosen to cover ± 1.5 standard deviations in player value, corresponding to roughly 90% confidence under normal uncertainty; values below 1.0 are too conservative, while $\Gamma > 2.0$ yields solutions indistinguishable from the nominal):

Table 12: Robust vs. Nominal Roster Optimization Comparison

Metric	Nominal MILP	Robust ($\Gamma = 1.5$)
Objective value (multi-year EV)	242.1	231.4
Uncertainty penalty	0	10.7
Average age	26.0	26.0
Total social followers	13.3M	13.3M
Roster composition	Max nominal EV	Same 12 players; penalty 4.4%

Interpretation: The robust formulation sacrifices 4.4% of nominal multi-year value (uncertainty penalty 10.7) to hedge against injury and projection uncertainty. On the 2024 data run, the nominal and robust rosters coincide (same 12 players, avg age 26.0, 13.3M social followers); the penalty is the cost of the hedge. With different player pools or parameters, the robust solution can tilt toward younger, lower-uncertainty players. In practice, a team with strong injury data or a win-now mandate might choose a lower Γ (closer to nominal); a team prioritizing long-term value might retain $\Gamma = 1.5$ or increase it.

11.3 Free Agency Discrete Choice: Computational Results

As in Section 10.3, coefficients are calibrated to match observed 2023–2024 free agency moves [5]. The resulting coefficients:

Table 13: Free Agency Utility Coefficients (Calibrated to 2023–2024 Moves)

Factor	Coefficient	Interpretation
β_{salary}	0.80	Log-salary effect (diminishing returns)
$\beta_{\text{championship}}$	1.20	Championship probability (strong weight)
β_{market}	0.60	Endorsement/market size
β_{tax}	0.40	After-tax income factor
$\beta_{\text{incumbent}}$	0.50	Bonus for staying with current team

Model Validation: The model correctly predicts the direction of observed moves:

- Skylar Diggins-Smith (Phoenix → Seattle): Model gives 24.5% to Seattle, consistent with tax-free state + improved championship odds.
- Nneka Ogumwike (Los Angeles → Seattle): Model gives 26.0% to Seattle, capturing tax + championship effects.

Hypothetical: Caitlin Clark Free Agency. If Clark reaches unrestricted free agency with supermax offers (\$249,244) from multiple teams, the model predicts:

Table 14: Caitlin Clark Free Agency Probabilities (Hypothetical)

Team	Tax Rate	Chip Prob	Signing Prob
New York Liberty	10.9%	25%	36.7%
Indiana Fever (incumbent)	3.2%	5%	34.2%
Las Vegas Aces	0%	22%	29.1%

Toronto Tax Problem: Toronto’s 53% marginal rate [23] means it must offer ~\$530K to match Las Vegas’s after-tax supermax (\$249K)—a structural free-agency disadvantage for high-tax expansion markets.

11.4 Superstar Index: Social Media Integration and Limitations

We augment the baseline SI (PER + minutes proxy) with Instagram follower data [17]:

$$SI^+ = 0.5 \cdot SI_{\text{base}} + 0.3 \cdot \frac{\text{Followers}_i}{\max_j \text{Followers}_j} + 0.2 \cdot \frac{\text{Engagement}_i}{\max_j \text{Engagement}_j}$$

where Engagement is approximated by follower growth rate (available for top players). For players without social data, we use SI_{base} alone. **Limitation:** A richer SI would incorporate real-time sentiment analysis (e.g., Twitter/X mentions, highlight-clip views, merchandise sales), including positive vs. negative sentiment (e.g., polarizing players) to refine revenue contribution. Such data require proprietary access; we use publicly available follower counts as the best available proxy.

Table 15: Top WNBA Players by Instagram Followers (Dec 2024) [17]

Player	Followers (M)	Note
Angel Reese	4.9	Exceeds WNBA official (3.2M)
Caitlin Clark	3.5	2.7M reach in one campaign
Sabrina Ionescu	1.5	Established All-Star
A’ja Wilson	≥ 1.0	2× MVP

Insight: WNBA players have more followers than their teams (no franchise exceeds 688K; 10+ players exceed 1M). Player brand drives franchise value more directly than in team-centric leagues.

11.5 Data and Assumptions Summary

All numerical inputs in this report are either (a) cited from a single public source (Basketball-Reference, Spotrac, Sportico, Forbes, WNBA, AP, ESPN, USA Today, LA Times, Tax Foundation), or (b) explicitly stated model parameters (e.g., $\rho(a)$, discount rate $r = 0.1$, $\Gamma = 1.5$, logit β s). No fabricated or unverifiable data are used. Model outputs (e.g., optimal roster, 3-year EV, signing probabilities) are computed from those inputs and parameters.

12 Conclusions

We developed an integrated framework (MILP + system dynamics + game theory + pricing) to maximize WNBA franchise profit and value. Key contributions:

1. **MRP (performance + brand).** By incorporating superstar externalities, our player valuation captures revenue effects missed by traditional Win Shares. The Clark effect (26.5% of league revenue) motivates weighting SI alongside Win Shares in roster and protection decisions.
2. **Optimal roster under constraints.** The MILP produces actionable rosters that balance performance and brand, respecting salary cap and positional rules. On 2024 data, the optimal roster uses 99.99% of cap and includes both high-WS stars (Wilson, Stewart) and cost-effective high-SI players (Clark, Boston).
3. **Multi-year financial planning.** System dynamics enables scenario analysis (expansion, injury, CBA changes) with explicit feedback loops. Variable agent fees (4–6% by star power) and capacity-constrained ticket pricing are integrated so that financial projections reflect realistic revenue and cost structure.
4. **Strategic protection.** Minimax ensures worst-case loss is minimized in expansion drafts. For Indiana, protecting the top six by value (Boston, Clark, Mitchell, Smith, Hull, Fagbenle) limits the loss to the seventh-best player (Wallace, value 1.00); exposing high-salary, low-value veterans (e.g., Wheeler at \$222k) can provide cap relief if selected.
5. **Revenue management.** Segment-based pricing with capacity constraints yields shadow prices for sell-out games (\$37 premium) and volume-building prices for weeknight games (\$11); the approach increases ticket revenue by 15–17% versus flat pricing.

Further refinements. Possible extensions include: reinforcement learning for real-time ticket pricing; play-by-play or social sentiment in SI; and a two-stage stochastic program that embeds the free-agency logit into roster optimization. The models in this report use only publicly available, cited data and stated assumptions. In practice, the framework is intended as a living dashboard: update player values and attendance data as the season progresses, re-run the MILP and protection analysis when the CBA is finalized, and adjust ticket pricing by segment as demand signals arrive.

12.1 Strengths and Limitations Summary

Table 16: Model Strengths and Limitations

Strengths	Limitations
MRP integrates performance and brand	SI relies on proxies when social data unavailable
MILP has guaranteed optimality	Free-agency behavior modeled separately (logit)
Closed-form ticket pricing; capacity constraint modeled	Isoelastic demand may not hold at extremes
Minimax has provable guarantee	Assumes expansion team value-maximizes
System dynamics with variable agent fees (4–6%)	Agent negotiation simplified; parameters from limited data
All data from public, verifiable sources	Private revenue data unavailable

Letter to the Owner and General Manager

Dear Owner and General Manager,

We have developed a dynamic decision-making model to support both competitive success and financial health for the 2026 season and beyond. Below we summarize our recommended strategy, trade-offs, and next steps.

Recommended Strategy.

1. **Roster Construction:** Prioritize players who drive both wins and revenue. Our analysis identifies Wilson and Stewart as high-value anchors, complemented by cost-effective rising stars like Clark and Boston; this balanced approach uses 99.99% of the \$1.5M cap. If the proposed \$5M cap is ratified, you will have flexibility to add 2–3 additional All-Stars.
2. **Expansion Draft:** The 2026 expansion exposes you to asset loss. We recommend a defensive protection strategy that shields your six most valuable players—those who contribute most to wins and brand equity. For Indiana, protect Clark, Boston, and Mitchell at all costs. Exposing a high-salary veteran with declining performance can provide cap relief if selected.
3. **Financial Planning:** The 2026 media deal raises league revenue to \$200M, but expansion dilutes per-team share to \$13.3M (from \$15M) [1]. Maintain cash reserves of at least \$2M to weather injury or attendance shocks. Investing in superstars may reduce short-term cash but raises franchise value over 3–5 years—WNBA valuations grew 180% in the last year alone [19].
4. **Ticket Pricing:** Segment your 20-game home schedule by opponent strength and day of week. For marquee matchups (e.g., hosting the Fever), premium pricing at \$37 per seat maximizes sell-out revenue; for weeknight games against weaker opponents, price at \$11 to build fan base. This segmented approach can increase total ticket revenue by 15–17% versus flat pricing [28].

Trade-Offs and Risks. Prioritizing revenue occasionally conflicts with pure competitive performance (a popular but average player may outrank a better but less marketable one). CBA uncertainty: if the cap stays at \$1.5M, roster flexibility is constrained; at \$5M, league-wide competition for superstars intensifies. Player brand metrics rely on public proxies; internal merchandise and social data would improve accuracy. Free agency, trades, and injuries add execution risk; the model can scenario-test these shocks but cannot eliminate them.

In closing, the framework aligns competitive success with financial returns. WNBA valuations have grown 180% [19]; building winning, marketable rosters drives long-term value. We recommend using this as a decision dashboard: update player values as the season unfolds, re-optimize when CBA terms are finalized, and adjust ticket pricing based on demand signals.

Sincerely,
The Modeling Team

References

- [1] Associated Press. WNBA announces landmark 11-year media rights deal with Disney, Amazon Prime and NBC. AP News, 2024. <https://apnews.com/article/wnba-media-deal-4d1a47fc4e0ee95d413322a96f8a7ce1>.
- [2] Basketball Reference. 2024 WNBA player stats – advanced. Basketball-Reference.com, 2024. https://www.basketball-reference.com/wnba/years/2024_advanced.html.
- [3] Dimitris Bertsimas, David B. Brown, and Constantine Caramanis. Theory and applications of robust optimization. *SIAM Review*, 53(3):464–501, 2011.
- [4] Joris Drayer, Daniel A. Rascher, and Chad D. McEvoy. Dynamic ticket pricing in sport: An agenda for research and practice. *Sport Marketing Quarterly*, 21(3):184–194, 2012.
- [5] ESPN. WNBA free agency and trade tracker 2024. ESPN, 2024. https://www.espn.com/wnba/story/_/id/39378597/2024-wnba-free-agency-trade-tracker.
- [6] ESPN. WNBA injuries: 2024 season tracker. ESPN, 2024. Cameron Brink ACL, Angel Reese wrist, Diamond Miller knee. <https://www.espn.com/wnba/injuries>.
- [7] Drew Fudenberg and Jean Tirole. *Game Theory*. MIT Press, 1991.
- [8] Andrew Garfinkel. The use of price-to-revenue ratios in valuing sports franchises. Technical Report 3421626, SSRN, 2019. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3421626.
- [9] Google. Gemini 3 pro. Large language model; used as described in Report on Use of AI Tools, 2026.
- [10] Jahn K. Hakes and Raymond D. Sauer. An economic evaluation of the Moneyball hypothesis. *Journal of Economic Perspectives*, 20(3):173–185, 2006.
- [11] Hans Kellerer, Ulrich Pferschy, and David Pisinger. *Knapsack Problems*. Springer, 2004.
- [12] Brett Knight. The WNBA’s most valuable teams 2025. Forbes, 2025. <https://www.forbes.com/sites/brett-knight/2025/06/06/wnba-valuations-2025/>.
- [13] Los Angeles Times Staff. Caitlin Clark responsible for 26.5% of WNBA revenue, expert says. Los Angeles Times, 2024. Analysis by Ryan Brewer, Indiana University. <https://www.latimes.com/sports/story/2024-12-04/caitlin-clark-worth-wnba-revenue>.
- [14] Daniel McFadden. Conditional logit analysis of qualitative choice behavior. *Frontiers in Econometrics*, pages 105–142, 1974.
- [15] Sherwin Rosen. The economics of superstars. *American Economic Review*, 71(5):845–858, 1981. JSTOR stable/1803469.
- [16] Gerald W. Scully. Pay and performance in major league baseball. *American Economic Review*, 64(6):915–930, 1974.
- [17] Sportico Staff. WNBA star popularity is unlike other u.s. sports leagues. Sportico, 2024. Instagram follower analysis. <https://www.sportico.com/leagues/basketball/2025/wnba-stars-instagram-1234852560/>.
- [18] Sportico Staff. WNBA team values 2024: Las Vegas, Seattle lead, average hits \$96M. Sportico, 2024. <https://www.sportico.com/valuations/teams/2024/wnba-team-values-1234784566/>.

- [19] Sportico Staff. WNBA team values 2025: Valkyries no. 1 at \$500M, average up 180%. Sportico, 2025. <https://www.sportico.com/valuations/teams/2025/wnba-team-values-1234857652/>.
- [20] Sports Business Journal. Portland awarded WNBA expansion franchise for record-setting price. Sports Business Journal, 2024. <https://www.sportsbusinessjournal.com/Articles/2024/09/18/portland-w-nba-expansion-franchise-announced-play-2026/>.
- [21] Spotrac. WNBA team salary cap tracker. Spotrac, 2025. <https://www.spotrac.com/wnba/cap/>.
- [22] John D. Sterman. *Business Dynamics: Systems Thinking and Modeling for a Complex World*. McGraw-Hill, 2000.
- [23] Tax Foundation. State individual income tax rates and brackets, 2024. Tax Foundation, 2024. <https://taxfoundation.org/data/all/state/2024/>.
- [24] Kenneth E. Train. *Discrete Choice Methods with Simulation*. Cambridge University Press, 2nd edition, 2009.
- [25] USA Today Sports. Explaining WNBA–WNBPA CBA negotiations: Issues, offers and what’s next. USA Today, 2025. <https://www.usatoday.com/story/sports/wnba/2025/12/03/87566711007/>.
- [26] Mechelle Voepel. 2026 WNBA expansion draft: Protected player predictions. ESPN, 2026. https://www.espn.com/wnba/story/_/id/47272305.
- [27] WNBA Communications. WNBA delivers record-setting 2024 season. WNBA.com, 2024. <https://www.wnba.com/news/wnba-delivers-record-setting-2024-season>.
- [28] Xiaomeng Xu, Serguei Nettekssine, and Pavan Sainam. Designing and evaluating dynamic pricing policies for major league baseball tickets. Technical Report 2796407, SSRN, 2016. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2796407.

Report on Use of AI Tools

Our team used **Google Gemini 3 Pro** (2026 version) to assist with (1) refining and choosing modeling procedures, (2) finding references and learning about theories/data sources, (3) wording and clarity refinement on draft paragraphs (all documented below), and **GitHub Copilot** (2026 version) for code auto-completions when writing the Python models. All critical thinking, model selection, and final conclusions were performed by the team. Below we document all uses: modeling (1–4), references/theories (5), wording (6), and code (7). Inline citations in the main report (e.g., [9] in §3.2, §3.4, §3.3, §3.5) indicate where each procedure was consulted.

1. Roster optimization: MILP vs alternatives

Query:

We are building a WNBA roster optimization model under a salary cap. We want to maximize both on-court performance (win shares) and off-court value (brand/superstar appeal). We are considering: (a) mixed-integer linear programming (MILP) with a composite objective, (b) a greedy heuristic, or (c) a simple ranking rule. What are the main trade-offs? Which approach best fits a 25-page contest report where we need to show math and run code?

Output:

MILP is well-suited when you have a linear objective (e.g., weighted sum of win shares and brand metric), linear constraints (cap, roster size, positions), and binary selection variables. It gives a provably optimal solution for that formulation and is standard in operations research, so judges will recognize it. Greedy heuristics are easier to code but do not guarantee optimality; ranking rules are transparent but ignore constraint interactions. Recommendation: use MILP (e.g., PuLP or similar) with a single composite objective so you can show the formulation, dual/pricing interpretation, and one clear optimal roster. If you want to hedge uncertainty in player values, you can add a robust optimization variant (e.g., budget of uncertainty) as an extension. You could also run a greedy heuristic as a benchmark to report solution quality; some judges like seeing a comparison.

We adopted MILP and the robust extension ($\Gamma = 1.5$); we did not implement the suggested greedy benchmark. The team chose PuLP, λ_1 , λ_2 , and all data sources ourselves.

2. Expansion draft: game-theoretic protection

Query:

We need to model an expansion draft where our WNBA team protects 6 players and an expansion team picks one from the rest. We want to minimize worst-case loss (value of the player we lose). Should we use minimax, or is there a simpler rule like “protect the 6 highest salary” or “protect the 6 highest win shares”? How do we formulate it for a contest report?

Output:

Treat the expansion team as an adversary who picks the player that hurts you most. Then your problem is: choose a protection set S of size k to minimize the maximum value of any player the adversary can take. That is minimax: $\min_S \max_{p \notin S} v_p$. The solution is to protect the k players with highest value v_p ; the worst-case loss is the value of the $(k + 1)$ th-best player. “Protect highest salary” or “highest win shares” can be suboptimal if value is a combination of performance and brand (e.g., a high-salary, low-value veteran is better to expose than a cost-effective star). Formulate v_p as your MRP or composite score so the protection set aligns with roster value. If you want to go further, you could model the expansion team’s selection as a probabilistic choice (e.g., they prefer high-value players with some noise) and optimize expected loss instead of worst-case; that would require more assumptions.

We used the minimax formulation and the MRP-based v_p ; we did not implement the probabilistic/expected-loss variant. We cited Fudenberg-Tirole for Stackelberg ourselves. The team chose λ_1 , λ_2 and applied the protection list to Indiana Fever.

3. Financial projection: system dynamics vs single-year

Query:

We want to project WNBA team finances over 5 years with brand equity, fan base, and cash. Should we use a multi-year system dynamics (stock-flow) model or a single-year snapshot? What state variables and flow equations would you suggest for a contest report?

Output:

A multi-year system dynamics model fits the “dynamic decision-making” prompt better than a static snapshot. Use state variables: brand equity B_t , fan base F_t , cash reserves C_t . Flows: brand grows with wins and star power, decays slightly each period; fan base grows with wins and marketing spend, decays otherwise; cash = revenue minus salaries, agent fees, operating costs. Revenue can depend on F_t and B_t (e.g., $R_t = \rho F_t(1 + \eta B_t)$). Use discrete time $t = 0, \dots, T$ (e.g., annual). Name parameters (α , δ , γ , etc.) so the reader sees how wins and star power enter. Run scenarios (e.g., high vs low CBA cap) to show sensitivity. You could also add a fourth state (e.g., arena quality or sponsorship pipeline) if you want more granularity.

We used the three-state structure, parameter naming, and scenario idea; we did not add a fourth state. We added variable agent fees (4–6% by star power) ourselves. The team calibrated all parameters and ran scenarios.

4. Ticket pricing: elasticity-based vs flat

Query:

We need to model ticket pricing for a WNBA team. We have different game segments (premium opponent, weeknight weak). Should we use a single price or segment-specific? What functional form (linear demand, isoelastic) is tractable and defensible for a contest?

Output:

Segment-specific pricing is standard in sports: premium games (marquee opponent, weekend) have less elastic demand than weeknight weak games. Use constant-elasticity demand $q = q_0(p/p_0)^{-\epsilon}$ per segment; then revenue $R = p \cdot q$ is maximized at $p^* = c\epsilon/(\epsilon - 1)$ when $\epsilon > 1$, where c is marginal cost. So you get a closed-form p^* per segment. If demand at p^* exceeds arena capacity, cap quantity at capacity and use the shadow price (price at which demand equals capacity). This gives you a clear table: segment, ϵ , p^* , q^* , revenue. Cite dynamic pricing literature (e.g., SSRN baseball ticket pricing) for credibility. Alternatively, you could use a linear demand $q = a - bp$ per segment; the math is simpler but the optimal price is less interpretable than the elasticity form.

We used the isoelastic formulation and capacity/shadow-price approach; we did not use linear demand. We cited Drayer and the SSRN paper. The team chose elasticity values and base prices from the literature and our own calibration.

5. Finding references and learning about theories/data

We used Gemini to identify or validate where to learn more about theories and where to obtain data. Below we document each use.

Use (a) — Marginal revenue product and roster optimization in sports.

Query:

Where can I learn more about marginal revenue product (MRP) in sports and using it for roster optimization? We need academic or credible sources for a math modeling report.

Output:

Classic references: Scully (1974) “Pay and Performance in Major League Baseball,” *American Economic Review*—foundational MRP in sports. For “Moneyball”-style use of performance metrics in roster decisions: Hakes and Sauer (2006) “An Economic Evaluation of the Moneyball Hypothesis,” *Journal of Economic Perspectives*. For superstar effects: Rosen (1981) “The Economics of Superstars,” *AER*. For optimization under a cap, operations research literature on knapsack and integer programming; PuLP or similar for implementation. Cite Scully for MRP decomposition, Rosen for superstar externalities, Hakes-Sauer for empirical performance metrics. You could also add Kahn (2000) on NBA salary structure or Rottenberg (1956) on the reserve clause for historical context if you want more breadth.

We cited Scully, Rosen, and Hakes-Sauer; we did not add Kahn or Rottenberg as they were not central to our formulation. We chose PuLP ourselves.

Use (b) — WNBA data sources: salaries, stats, valuations.

Query:

What are reliable public data sources for WNBA player salaries, advanced stats (win shares, PER), and team valuations? We need to build a dataset for a 2024–2025 modeling report.

Output:

Salaries and cap: Spotrac (wnba cap tracker) is the standard public source for NBA/WNBA salaries and cap. Advanced stats: Basketball-Reference has WNBA advanced stats (Win Shares, PER, minutes, etc.) by season. Valuations: Sportico and Forbes publish WNBA team valuations annually. Attendance: WNBA official releases. For a contest report, cite Spotrac for cap/salaries, Basketball-Reference for Win Shares and PER, Sportico/Forbes for valuations. Cross-check key numbers (e.g., cap \$1.5M) across sources. For social/media metrics (e.g., Instagram followers for your Superstar Index), Sportico and similar outlets publish WNBA star popularity pieces; you could also scrape public follower counts if needed.

We used Spotrac, Basketball-Reference, Sportico, and Forbes; we did not scrape follower counts and instead used our PER/minutes proxy and later Sportico social for a separate subsection. We use Spotrac only for the cap figure (single source). We cited these in Data and Metrics and Assumptions.

Use (c) — Dynamic ticket pricing in sports.

Query:

What are good references for dynamic ticket pricing in sports? We need an elasticity-based or demand-based formulation for a contest report, ideally with a citable paper.

Output:

Drayer et al. (2012) “Dynamic Ticket Pricing in Sport: An Agenda for Research and Practice,” *Sport Marketing Quarterly*—overview of segment-based and dynamic pricing in sport. For empirical implementation and elasticity: Xu, Netessine, Sainam (2016) SSRN “Designing and Evaluating Dynamic Pricing Policies for Major League Baseball Tickets”—they use demand elasticity by segment and revenue maximization. For discrete choice or multinomial logit in player/team decisions: Train (2009) “Discrete Choice Methods with Simulation,” Cambridge; McFadden (1974) for conditional logit. Cite Drayer for sport context, Xu et al. for elasticity and revenue formulation. There is also a growing literature on real-time or reinforcement-learning-based pricing (e.g., MDPs); you could mention that as future work if you don’t implement it.

We cited Drayer, the SSRN paper, Train, and McFadden. We did not add a citation for RL/MDP pricing (we mention it briefly as future work without a separate reference). We did not cite every possible reference the output could imply.

6. Wording and clarity refinement

We used Gemini for wording and clarity on draft paragraphs. We document **every** such use below (all listed).

Use (a) — Letter to the Owner (expansion draft).

Query:

Please suggest a clearer way to phrase this sentence for a technical report: “We recommend that management use our minimax protection model to choose the six players whose loss would be most damaging.” We want it to sound more business-focused and less jargon-heavy for a letter to the owner.

Output:

A more business-focused version: “We recommend a defensive protection strategy that shields your six most valuable players—those who contribute the most to both wins and brand equity.” This avoids the term “minimax” and emphasizes the outcome (protecting value) rather than the method. You can add: “For Indiana, this means protecting Clark, Boston, and Mitchell at all costs.” if the audience is team-specific.

We adopted the first sentence with minor wording changes; we added the Indiana-specific line (Clark, Boston, Mitchell) ourselves from our protection table.

Use (b) — Trade-offs and risks paragraph.

Query:

Shorten and tighten this bullet list for a letter. Keep the same four points: (1) wins vs profit conflict, (2) CBA uncertainty, (3) data limitations on brand value, (4) execution risk. Make it one short paragraph instead of four bullets.

Output:

Prioritizing revenue occasionally conflicts with pure competitive performance (a popular but average player may outrank a better but less marketable one). CBA uncertainty: if the cap stays at \$1.5M, roster flexibility is constrained; at \$5M, league-wide competition for superstars intensifies. Player brand metrics rely on public proxies; internal merchandise and social data would improve accuracy. Free agency, trades, and injuries add execution risk; the model can scenario-test these shocks but cannot eliminate them.

We used the output as the basis and tightened a few phrases (e.g., “scenario-test” and cap numbers) to match our report.

Use (c) — Robust optimization parameter (Γ).

Query:

We use robust optimization with a budget of uncertainty $\Gamma = 1.5$. We need one sentence for the report explaining why we chose 1.5 and not 1.0 or 2.0. Keep it brief and cite-worthy.

Output:

You could say: “ $\Gamma = 1.5$ is chosen to cover approximately ± 1.5 standard deviations in player value, corresponding to roughly 90% confidence under normal uncertainty; values below 1.0 are too conservative, while $\Gamma > 2.0$ yields solutions indistinguishable from the nominal formulation.” That justifies the choice without over-claiming.

We used the output sentence with a small edit (e.g., “cite-worthy” → our wording) and placed it in the Robust Roster Optimization Results subsection.

Use (d) — Closing paragraph of the Letter.

Query:

We have a closing paragraph that says the framework aligns wins with financial returns and that WNBA valuations grew 180%. We want to end with a sentence that tells the owner to use the model as a living tool. Suggest one sentence.

Output:

“We recommend using this as a decision dashboard: update player values as the season unfolds, re-optimize when CBA terms are finalized, and adjust ticket pricing based on demand signals.” That frames the model as actionable rather than one-time.

We used the dashboard sentence with minor edits (e.g., “re-optimize when CBA terms are finalized”) to match our Letter.

Use (e) — Summary: main conclusion lead.

Query:

We want the summary of our report to start with our main conclusion so the reader is hooked. We show MRP + minimax + segment pricing, 15–17% revenue gains, 180% valuations. Suggest one opening sentence that states this clearly.

Output:

“We show that valuing players by Marginal Revenue Product (performance + brand), using minimax protection in the 2026 expansion draft, and deploying segment-based dynamic ticket pricing aligns wins with profit and yields 15–17% revenue gains; WNBA valuations grew 180%, justifying investment in superstars despite short-term cost.” You can add a citation for the 15–17% and 180% claims.

We adapted the output for our Summary, added the citations ourselves, and split into two sentences for readability.

Use (f) — Introduction: mapping problem questions to our framework.

Query:

We have five problem questions (dynamic model, acquisition, expansion, ticket pricing, injury). We need one short paragraph that maps them to our solution structure (system dynamics, MILP, minimax, pricing, scenario re-runs) so the reader sees the roadmap.

Output:

“Our solution maps these to an integrated framework: (1) system dynamics for multi-year financial and brand evolution; (2) MILP for roster optimization under the cap; (3) minimax game theory for expansion-draft protection; (4) elasticity-based pricing for ticket revenue; and (5) scenario re-runs (MILP, dynamics, pricing) when injury or CBA changes occur. The remainder of the report develops each component and then combines them into a single decision support system.”

We used the mapping paragraph with one clause reordered for flow and placed it at the end of Problem Formulation.

Use (g) — Data Sources subsection: one-sentence description.

Query:

Shorten this: “Player statistics (Win Shares, PER, minutes) are taken from Basketball-Reference WNBA Advanced Stats and implemented in our data module; we do not use play-by-play APIs.” Keep it one sentence, same meaning.

Output:

“Player statistics (Win Shares, PER, minutes) are taken from Basketball-Reference [2] and implemented in our data module; we do not use play-by-play APIs.”

We had already cited Basketball-Reference; the output matched our intended sentence so we kept it and added the citation.

Use (h) — Sensitivity table: interpretation paragraph.

Query:

We have a sensitivity table (cap, λ_2 , injury, protection 5 vs 6, elasticity, media). We need one short paragraph interpreting the most important rows for the reader (cap, λ_2 , injury, protection).

Output:

“The \$1.5M \rightarrow \$5M cap change has the largest positive impact on objective value (+42%) and 5-year cash (+18%), confirming that CBA outcome is critical. Increasing λ_2 (brand weight) shifts the roster toward high-SI players and raises both Z^* and cash, while $\lambda_2 \rightarrow 0$ (pure performance) reduces both. The injury scenario ($\sigma_t \rightarrow 0.7\sigma_t$) shows a 22% drop in 5-year cash with a 2-year recovery, justifying the \$2M reserve. Protecting 5 instead of 6 players sharply increases worst-case loss, so the expansion draft protection decision is highly consequential.”

We used the output as the basis and added the protection (5 vs 6) interpretation ourselves from our robustness results.

Use (i) — Conclusions: living dashboard sentence.

Query:

We want to end the Conclusions section with one sentence saying the framework should be used as a living tool (update inputs, re-run when CBA finalizes, adjust pricing). Suggest one sentence.

Output:

“In practice, the framework is intended as a living dashboard: update player values and attendance data as the season progresses, re-run the MILP and protection analysis when the CBA is finalized, and adjust ticket pricing by segment as demand signals arrive.”

We integrated this sentence into the Further refinements paragraph with minor wording changes.

For each of the above (a)–(i), the team chose which suggestions to adopt, edited all text, and often combined AI output with our own phrasing or prior decisions. We did not use AI to generate substantive content or conclusions. The list above is the complete set of wording uses.

7. Code: GitHub Copilot

GitHub Copilot (2026 version). Auto-completions for code used in preparing our Python models (roster MILP, system dynamics, ticket pricing, visualizations). We did not use Copilot to generate full functions or algorithms; we used it for inline completions (e.g., variable names, common library calls). All logic and model implementation were written and verified by the team.

All model choices, implementations, and conclusions in the 25-page solution are the responsibility of the team.