System Requirements Specification

Career Strategy Intelligence System

Version 1.0

Technical Documentation

February 13, 2026

Contents

1 Introduction

1.1 Purpose

This System Requirements Specification defines the functional and non-functional requirements for a career strategy intelligence system designed to optimize long-term career trajectories for highly ambitious users through causal reasoning and strategic analysis under severe computational constraints.

The system addresses the fundamental problem of career decision-making as a strategic optimization problem under uncertainty, where candidates must maximize expected lifetime utility across multiple uncertain futures while operating under finite resources, incomplete information, and irreversible constraints.

1.2 Scope

1.2.1 In Scope

- Strategic career decision support through causal inference
- Market regime intelligence and temporal dynamics analysis
- Institutional and advisor quality assessment from public signals
- Counterfactual trajectory analysis for persist-versus-pivot decisions
- Negotiation leverage analysis and timing optimization
- User self-calibration and preference discovery under uncertainty
- Longitudinal outcome tracking and system learning

1.2.2 Out of Scope

- Job board functionality with comprehensive listings
- Social networking features and viral growth mechanisms
- Resume editing and document generation services
- Employer relationship management and partnerships
- Direct application submission infrastructure
- Human counseling or emotional support services
- Real-time messaging and conversational interfaces

1.3 Target User Profile

1.3.1 Primary Users

- PhD candidates and postdoctoral researchers
- Highly competitive job market participants in knowledge work

- Individuals facing high-stakes, irreversible career decisions
- Users capable of probabilistic reasoning and uncertainty tolerance
- Users willing to contribute longitudinal outcome data

1.3.2 User Characteristics

- Sophisticated career goals beyond immediate employment
- Values strategic positioning over application volume
- Comfortable with confidence intervals and probabilistic claims
- Seeks truth and calibrated advice over reassurance
- Willing to receive uncomfortable feedback on career trajectory
- Capable of executing multi-step strategic recommendations

1.4 System Overview

The system provides strategic intelligence through analysis of public data sources, user-contributed outcome data, causal inference models, regime-level market monitoring, and personalized prediction engines operating within strict computational constraints imposed by free-tier API access and consumer hardware limitations.

1.5 Document Conventions

Requirements are identified using the following nomenclature:

- **FR-X.Y.Z**: Functional Requirement (Section X, Subsection Y, Item Z)
- **NFR-X.Y.Z**: Non-Functional Requirement
- **CC-X.Y.Z**: Computational Constraint
- **A-X.Y**: Assumption (Section X, Item Y)
- **OBJ-X.Y**: System Objective

Priority levels: **Critical**, **High**, **Medium**, **Low**

2 Problem Analysis

2.1 Core Problem Statement

Job and PhD seeking constitutes a strategic optimization problem under uncertainty where candidates must maximize expected lifetime utility across multiple uncertain futures while constrained by finite resources (time, money, cognitive capacity), irreversible or costly-to-reverse decisions, incomplete information about opportunity quality, dynamic competitive landscapes, and non-stationary evaluation criteria.

The fundamental challenge: candidates must commit to decisions before knowing which path would have been optimal, operating in an environment where the evaluation criteria themselves co-evolve adversarially with candidate strategies.

2.2 Structural Challenges

2.2.1 Information Asymmetry

Employer and Program Advantage Institutions possess information candidates cannot access:

- True funding stability and internal political dynamics
- Actual day-to-day work culture versus marketed culture
- Historical retention and success rates not publicly disclosed
- Whether posted positions reflect genuine need or political theater
- For PhD programs: advisor temperament, lab group dynamics, project viability

This creates a winner's curse scenario where those most eager to accept positions may be precisely those who most severely misestimate the offering quality.

Candidate Private Information Candidates hold private information they cannot credibly prove:

- Actual capabilities versus credential proxies
- Resilience, creativity, and work ethic
- Risk tolerance and reservation utility thresholds
- Domain knowledge not reflected in formal metrics

The signaling trap emerges: candidates cannot simply declare their quality since cheap talk is discounted, creating massive inefficiency through pooling equilibria where high-quality candidates become indistinguishable from low-quality candidates.

2.2.2 Signaling Problems

Candidates must invest in costly signals that correlate with underlying quality:

- Publications for academic tracks
- Credentials from high-status institutions
- Recommendation letters from recognized authorities
- Project portfolios and employment at respected firms

This generates arms race dynamics where signal thresholds rise as more candidates invest, requiring candidates to run faster merely to maintain position. The burden exhibits a U-shaped pattern: mid-tier candidates face the highest signaling costs, while those at the very top can afford countersignaling through strategic restraint.

Conservative institutions optimize to minimize Type I errors (accepting bad candidates) at the expense of Type II errors (rejecting good candidates), imposing asymmetric costs on candidates who bear the consequences of institutional risk aversion.

2.2.3 Competitive Positioning

Success depends on relative rank within comparison sets, not absolute quality, creating:

- Reference group dependencies: success contingent on unknown, stochastic competitor pools
- Threshold effects: binary outcomes at ranking cutoffs (3rd versus 4th when three positions exist)
- Strategic timing dependencies: optimal application timing depends on others' timing
- Explore-exploit dilemmas: resource allocation across reach, target, and safety opportunities with non-divisible applications, opaque correlation structures, and shifting opportunity sets

2.2.4 Career Trajectory as Markov Decision Process

Current state (credentials, age, employment status) determines accessible future states, generating:

- Path dependencies: early choices constrain later possibilities in non-obvious ways
- Option value decay: time spent in certain trajectories reduces value for alternative paths
- Secretary problem variants: sequential opportunities observed with noise, accept-reject decisions without full recall
- Exhaustion traps: reservation utility drifts downward during extended search, increasing acceptance probability for marginal offers precisely when evaluation capability is compromised
- Hysteresis: switching thresholds exceed zero by unknown margins due to irreversibility costs

Many career transitions are approximately irreversible:

- Leaving PhD programs wastes sunk investment and damages signaling credibility
- Transitioning from industry to PhD after five-plus years faces age and skill-staleness penalties
- Geographic moves impose relationship and network costs

2.2.5 Fundamental Impossibility

Evaluation criteria co-evolve adversarially with candidate optimization strategies:

1. Candidates optimize for observable metrics (citation counts, prestige markers, specific skill sets)
2. Metrics become noisier signals of underlying quality through Goodhart's Law
3. Competition drives up investment costs for acquiring metrics
4. Metrics become less predictive of actual job performance
5. Evaluators recognize degradation and shift to new criteria
6. Cycle restarts with new target metrics

This creates an optimization problem on a fitness landscape that reshapes itself in response to optimization efforts, rendering long-term planning fundamentally uncertain.

2.2.6 Uncertainty Cascade

Decision-making involves stacked uncertainties that multiply rather than add:

1. Preference uncertainty: what outcomes the candidate actually wants
2. Market uncertainty: what opportunities exist now and will exist in future
3. Rank uncertainty: candidate's competitive position in applicant pools
4. Information uncertainty: true quality of specific opportunities
5. Criterion uncertainty: how to evaluate opportunities given unknown preferences
6. Growth uncertainty: how candidate's capabilities will develop over time

The multiplicative nature of these uncertainties produces confidence intervals so wide as to approach meaninglessness for long-term optimization, necessitating satisficing strategies rather than optimization.

2.3 Rational Response Framework

Given the structural impossibility of optimization, the rational approach involves:

- Satisficing rather than optimizing (accepting good-enough rather than searching for optimal)
- Developing robust heuristics that perform acceptably across multiple scenarios

- Maintaining optionality to preserve future flexibility
- Gathering information efficiently to reduce critical uncertainties
- Accepting navigation in fog rather than solving clean maximization problems

3 Explicit Assumptions

3.1 User Assumptions

3.1.1 Target User Capabilities

A-1.1 Users can interpret quantitative risk assessments and probabilistic claims

A-1.2 Users are comfortable reasoning with uncertainty and confidence intervals

A-1.3 Users can make decisions with incomplete information

A-1.4 Users are willing to receive uncomfortable truths about career prospects

A-1.5 Users can execute multi-step strategic recommendations over extended timelines

3.1.2 Target User Limitations

A-1.6 Users experience preference uncertainty and may not know what they truly want

A-1.7 Users are subject to cognitive biases including impostor syndrome and Dunning-Kruger effects

A-1.8 Users have limited time and financial resources for career search activities

A-1.9 Users have emotional investment in career decisions that may cloud judgment

A-1.10 Users may have poor calibration in self-assessment of competitive position

3.1.3 User Contribution Willingness

A-1.11 Users will voluntarily contribute longitudinal outcome data years after initial decisions

A-1.12 Users will provide truthful reporting of career outcomes for system learning

A-1.13 Users accept privacy-preserving aggregation of their data to benefit other users

A-1.14 Users are willing to invest time in detailed profile setup for higher-quality recommendations

3.2 Market Assumptions

3.2.1 Labor Market Structure

A-2.1 Opportunities arrive sequentially over time rather than simultaneously

A-2.2 Decision-makers optimize primarily for risk minimization rather than upside maximization

A-2.3 Prestige signals carry weight beyond their intrinsic informational value

A-2.4 Hidden information asymmetries are pervasive and systematic across institutions

A-2.5 Institutional constraints affect hiring decisions beyond stated criteria

A-2.6 Markets shift between buyer's markets and seller's markets with detectable patterns

A-2.7 Fields and institutions follow ascent-plateau-decline trajectories over multi-year timescales

3.2.2 Information Environment

A-2.8 Public signals exist that correlate with hidden institutional quality

A-2.9 Historical outcome data is obtainable through public records and user contributions

A-2.10 Patterns in institutional behavior are detectable through systematic analysis

A-2.11 Funding and policy documents are publicly accessible and contain forward-looking information

A-2.12 Publication and hiring data is partially available through academic databases

A-2.13 Some private information remains fundamentally unknowable regardless of analysis effort

3.3 Technical Assumptions

3.3.1 Computational Environment

A-3.1 System must operate on consumer hardware without GPU acceleration

A-3.2 Free-tier API access with low tokens-per-minute and requests-per-minute limits

A-3.3 Token usage must be minimized through reasoning compression and caching

A-3.4 Pre-computation and batch processing are viable strategies for regime-level analysis

A-3.5 Local caching can provide acceptable staleness trade-offs for most queries

3.3.2 Data Availability

A-3.6 Public data sources are accessible through web scraping within legal bounds

A-3.7 User-contributed outcome data will accumulate over years to enable causal inference

A-3.8 Longitudinal tracking is feasible with user consent and engagement

A-3.9 Causal inference is possible with sufficient data despite selection bias

A-3.10 Selection bias in outcome reporting can be modeled and partially corrected

3.3.3 Algorithmic Capabilities

A-3.11 Lightweight statistical models suffice for most analysis tasks

A-3.12 Large language model calls are expensive and should be minimized to 5-10% of operations

A-3.13 Structured data processing is orders of magnitude cheaper than text generation

A-3.14 Batch processing with monthly or quarterly updates is acceptable for regime-level analysis

A-3.15 Heuristics can be learned from outcome data and compressed for zero-cost application

3.4 Product Assumptions

3.4.1 Value Proposition

A-4.1 Users value strategic intelligence over comprehensive information volume

A-4.2 Truth-telling is more valuable than engagement optimization for target users

A-4.3 Long-term trajectory optimization justifies subscription pricing

A-4.4 Computational honesty differentiates from engagement-optimized competitors

A-4.5 Users will accept and act on stop-searching advice when warranted

3.4.2 Competitive Environment

A-4.6 Existing platforms have structural incentive misalignments with user outcomes

A-4.7 Competitors cannot pivot to truth-telling without fundamental revenue model changes

A-4.8 Causal inference capability creates sustainable technical moat

A-4.9 First-mover advantage exists in outcome data accumulation

A-4.10 Network effects are achievable through user outcome contributions

4 Derived System Objectives

4.1 Primary Objectives

4.1.1 User Outcome Optimization

OBJ-1.1 Maximize expected lifetime career utility across uncertain futures

OBJ-1.2 Minimize wasted effort on low-probability opportunities

OBJ-1.3 Preserve option value across sequential decision points

OBJ-1.4 Enable informed decision-making under irreducible uncertainty

OBJ-1.5 Reduce career-altering mistakes through early warning systems

4.1.2 Strategic Intelligence Provision

OBJ-1.6 Detect market regime shifts three to five years before they become obvious

OBJ-1.7 Identify field and institutional trajectories (ascent, plateau, decline)

OBJ-1.8 Predict advisor and manager quality from public signals

OBJ-1.9 Calculate opportunity costs and trade-offs explicitly

OBJ-1.10 Model counterfactual career paths for persist-versus-pivot decisions

4.1.3 Computational Efficiency

OBJ-1.11 Minimize token consumption through reasoning compression

OBJ-1.12 Avoid unnecessary agent recursion and verbose outputs

OBJ-1.13 Prefer structured reasoning over chain-of-thought expansion

OBJ-1.14 Maximize use of cached and pre-computed analyses

OBJ-1.15 Operate within free-tier API constraints without degradation

4.2 Secondary Objectives

4.2.1 Calibration and Learning

OBJ-2.1 Improve predictions through continuous outcome feedback

OBJ-2.2 Detect systematic biases in reasoning components

OBJ-2.3 Update priors based on observed reality versus predictions

OBJ-2.4 Learn user-specific patterns and comparative advantages

OBJ-2.5 Compress successful reasoning patterns into reusable heuristics

4.2.2 User Self-Knowledge

OBJ-2.6 Disambiguate impostor syndrome from genuine skill gaps

OBJ-2.7 Identify comparative advantages relative to realistic comparison sets

OBJ-2.8 Detect preference uncertainty and recommend exploration strategies

OBJ-2.9 Project future preference changes for regret minimization

OBJ-2.10 Correct miscalibration in user self-assessment

4.2.3 Competitive Differentiation

OBJ-2.11 Provide capabilities structurally impossible for competitors

OBJ-2.12 Maintain incentive alignment prioritizing user outcomes over engagement

OBJ-2.13 Build defensible data moats through unique datasets

OBJ-2.14 Resist institutional capture and maintain truth-telling independence

OBJ-2.15 Create network effects through user-contributed outcome data

5 Functional Requirements

5.1 Temporal Strategy Capabilities

5.1.1 Path Irreversibility Analysis

Requirement ID: FR-1.1.1

Priority: High

Description: The system shall quantify future flexibility preserved or destroyed by accepting specific career opportunities.

Inputs:

- Opportunity details: position type, institution, field, geographic location, constraints
- User profile: current credentials, career stage, stated goals, constraints
- Current career trajectory: past positions, decisions, outcomes

Outputs:

- Option value quantification in terms of accessible future opportunity sets
- Time until path reversal becomes viable measured in years
- Accessible future opportunity sets from each choice with probability distributions
- Confidence intervals on projections at five-year and ten-year horizons

Acceptance Criteria:

- System identifies irreversible versus reversible decisions
- System estimates reversal costs including time, credential depreciation, and opportunity costs
- System provides separate five-year projections (higher confidence) and ten-year projections (lower confidence)
- System quantifies which future paths remain accessible versus foreclosed

Derived From: OBJ-1.3, A-2.7

Traces To: Agent capability for modeling state-dependent opportunity sets

5.1.2 Search Termination Optimization

Requirement ID: FR-1.1.2

Priority: High

Description: The system shall determine optimal search termination given user's prior distribution over opportunity quality and remaining time horizon.

Inputs:

- Application history: submissions, outcomes, timing of responses
- Remaining search time available before forced decision
- Current best offer if any exists
- User's reservation quality threshold

Outputs:

- Stop or continue recommendation with justification
- Updated reservation threshold based on Bayesian updating
- Expected value of continued search accounting for costs
- Probability distribution over future outcomes conditional on continuation

Acceptance Criteria:

- System uses Bayesian updating as applications resolve
- System accounts for search costs: financial, temporal, psychological
- System adapts threshold dynamically as information arrives
- System provides explicit stopping rules with confidence levels
- System can recommend stopping even when additional opportunities exist

Derived From: OBJ-1.2, A-2.1

Traces To: Agent capability for secretary problem variants with Bayesian updating

5.1.3 Decision Sequencing Optimization

Requirement ID: FR-1.1.3

Priority: Medium

Description: The system shall determine optimal order for decision-making given deadline topology and information revelation schedule.

Inputs:

- Multiple opportunities with associated deadlines
- Decision dependencies specifying which choices foreclose others
- Expected information arrival times for when decisions will resolve

Outputs:

- Optimal decision sequence maximizing option value
- Conditional acceptance strategies for opportunities with dependencies
- Identification of forced choices where opportunities are mutually exclusive

- Timeline visualization showing critical path and decision points

Acceptance Criteria:

- System identifies critical path applications that must be decided first
- System detects which applications become irrelevant if earlier decisions resolve favorably
- System prevents wasted effort on dominated opportunities
- System handles temporal constraints where accepting both A and B is impossible
- System accounts for information revelation timing in sequencing recommendations

Derived From: OBJ-1.3, A-2.1

Traces To: Agent capability for constraint satisfaction and graph analysis

5.2 Information Arbitrage Capabilities

5.2.1 Field Collapse Detection

Requirement ID: FR-1.2.1

Priority: High

Description: The system shall detect research areas or job categories entering terminal decline three to five years before decline becomes obvious to participants.

Inputs:

- Field or subfield identifier
- Historical metrics: funding trends, hiring patterns, publication metrics, enrollment data

Outputs:

- Decline probability with confidence interval
- Timeline projection indicating when collapse will become obvious
- Leading indicators that triggered detection
- Adjacent field pivot recommendations with credential transferability analysis

Acceptance Criteria:

- System monitors funding trends including agency budgets and grant success rates
- System tracks hiring patterns including new positions and time-to-fill metrics
- System analyzes publication metrics including venue prestige and citation patterns
- System detects enrollment trends including PhD applications and graduation rates
- System updates annually with new data

- System provides separate three-year and five-year projections with different confidence levels
- System achieves advance warning of at least three years before decline becomes consensus view

Derived From: OBJ-1.6, OBJ-1.5

Traces To: Agent capability for longitudinal trend analysis and regime detection

5.2.2 Funding Regime Shift Prediction

Requirement ID: FR-1.2.2

Priority: Medium

Description: The system shall predict changes in funding priorities twelve to twenty-four months before they become widely recognized by field participants.

Inputs:

- Field or subfield specification
- Agency strategic plans and policy documents
- Congressional testimony and budget proposals
- Program officer movements and promotions

Outputs:

- Funding trajectory classification: increasing, stable, or declining
- Emerging priority areas with confidence assessments
- Strategic alignment recommendations for research agenda
- Confidence levels on predictions at twelve-month and twenty-four-month horizons

Acceptance Criteria:

- System analyzes policy documents for priority shifts
- System tracks which program officers are promoted or hired as legitimacy signals
- System monitors conference session proposals for emerging theme detection
- System detects when established researchers pivot toward new topics
- System updates quarterly with new data
- System provides separate twelve-month and twenty-four-month projections
- System achieves advance warning of at least twelve months before shifts become consensus

Derived From: OBJ-1.6, A-2.11

Traces To: Agent capability for policy analysis and forward-looking intelligence

5.2.3 Advisor and Manager Quality Prediction

Requirement ID: FR-1.2.3

Priority: High

Description: The system shall estimate probability that a specific advisor or manager will be professionally destructive based solely on public signals to avoid legal liability.

Inputs:

- Advisor or manager identifier: name, institution, department
- Public record data accessible without authentication

Outputs:

- Risk level classification: Low, Moderate, or High
- Observable pattern analysis with specific evidence
- Diagnostic interview questions for user to ask during evaluation
- Comparison to field norms for contextual interpretation

Acceptance Criteria:

- System analyzes time-to-degree distributions where abnormally long indicates red flag
- System examines authorship patterns comparing advisor-first versus student-first papers
- System calculates turnover rates inferred from lab website changes over time
- System detects publication gaps where sudden silence may indicate crisis
- System uses public data only to avoid defamation risk
- System provides field-normalized comparisons not absolute assessments
- System generates specific diagnostic questions for user interviews
- System maintains fact-based analysis without speculation

Derived From: OBJ-1.5, A-2.8

Traces To: Agent capability for pattern detection from public signals

5.3 Strategic Positioning Capabilities

5.3.1 Prestige Arbitrage Identification

Requirement ID: FR-1.3.1

Priority: Medium

Description: The system shall identify when institutional prestige is mispriced relative to actual opportunity quality.

Inputs:

- Multiple opportunity comparisons with different prestige levels
- Institutional prestige indicators from rankings and reputation
- Outcome metrics including placement rates, funding success, and satisfaction

Outputs:

- True value gap versus prestige gap quantification
- Arbitrage opportunity identification with expected value
- Expected value comparison accounting for opportunity costs
- Recommendation with explicit justification for deviation from prestige ranking

Acceptance Criteria:

- System compares prestige tier to outcome metrics
- System identifies systematic deviations indicating overrated or underrated institutions
- System accounts for opportunity costs including salary, location, and stress
- System provides field-specific analysis not generic rankings
- System uses only public outcome data for defensibility
- System quantifies prestige premium in terms of years of career advantage

Derived From: OBJ-2.11, A-2.3

Traces To: Agent capability for comparing stated versus revealed institutional quality

5.3.2 Institutional Trajectory Forecasting

Requirement ID: FR-1.3.2

Priority: Medium

Description: The system shall predict whether an institution or department is ascending, stable, or declining over five to ten year horizons.

Inputs:

- Institution and department identifiers
- Historical performance data over multiple years

Outputs:

- Trajectory classification: Ascending, Stable, or Declining
- Growth or decline indicators with quantitative evidence
- Risk factors that could alter trajectory
- Timeline with confidence intervals at five-year and ten-year horizons

- Sustainability assessment of current trajectory

Acceptance Criteria:

- System tracks faculty and employee hiring and retention rates
- System monitors student and customer demand trends
- System analyzes research expenditure trajectories
- System evaluates financial health indicators
- System assesses strategic plan realism versus fantasy
- System provides separate five-year (higher confidence) and ten-year (lower confidence) projections
- System identifies whether trajectory is driven by sustainable or temporary factors

Derived From: OBJ-1.7, A-2.7

Traces To: Agent capability for time-series analysis and trajectory modeling

5.3.3 Market Regime Detection

Requirement ID: FR-1.3.3

Priority: High

Description: The system shall identify when macro labor market conditions have shifted requiring fundamental changes to search strategy.

Inputs:

- Field and career stage specification
- User's search history and outcomes
- Aggregate market indicators across the field

Outputs:

- Current regime classification: Buyer's market, Seller's market, or Transitional
- Regime-appropriate strategy recommendations
- Expected regime duration with confidence interval
- Indicators being monitored for regime change detection

Acceptance Criteria:

- System monitors job posting volumes over time
- System tracks applications-per-posting ratios
- System analyzes time-to-decision trends

- System calculates offer-to-application ratios
- System detects regime transitions with statistical significance
- System updates monthly with new data
- System provides regime-specific tactical guidance
- System distinguishes temporary fluctuations from persistent regime shifts

Derived From: OBJ-1.6, A-2.6

Traces To: Agent capability for regime detection and strategy adaptation

5.4 Counterfactual Analysis Capabilities

5.4.1 Career Path Simulation

Requirement ID: FR-1.4.1

Priority: High

Description: The system shall model what would have happened if user had made different past career choices using causal inference.

Inputs:

- User's actual career trajectory with decisions and outcomes
- Past decision points with alternatives considered
- Alternative choices not taken at each decision point

Outputs:

- Comparison of actual versus counterfactual outcomes
- Statistical confidence on counterfactual estimates
- Persist versus pivot recommendation based on evidence
- Matched cohort identification showing similar users who made different choices

Acceptance Criteria:

- System uses causal inference methods not mere correlation
- System identifies synthetic controls meaning similar users who made different choices
- System controls for selection effects in outcome comparisons
- System provides confidence intervals on counterfactual estimates
- System addresses should-I-have questions with empirical evidence
- System distinguishes causal effects from selection effects

Derived From: OBJ-1.10, A-3.9

Traces To: Agent capability for causal inference and synthetic control construction

5.4.2 Information Gap Prioritization

Requirement ID: FR-1.4.2

Priority: Medium

Description: The system shall generate diagnostic questions that maximize expected value of information for specific opportunities.

Inputs:

- Opportunity details with known and unknown attributes
- Current uncertainty set enumerating critical unknowns
- User's decision criteria and priorities

Outputs:

- Prioritized question list ordered by expected information value
- Expected information value per question quantified
- Question-assignment recommendations specifying who to ask what
- Evasion detection guidance for interpreting non-answers

Acceptance Criteria:

- System enumerates critical unknowns blocking good decisions
- System designs questions that partition hypothesis space efficiently
- System suggests question ordering specifying which to ask first
- System identifies evasive answers as informative signals
- System distinguishes resolvable versus irreducible uncertainty
- System calculates expected value of information for prioritization

Derived From: OBJ-1.4, A-2.13

Traces To: Agent capability for information-theoretic question design

5.4.3 Negotiation Leverage Analysis

Requirement ID: FR-1.4.3

Priority: Medium

Description: The system shall quantify user's BATNA (Best Alternative To Negotiated Agreement) and determine optimal negotiation strategy for specific offers.

Inputs:

- Offer details including salary, startup funds, benefits, and constraints

- User's alternatives including other offers and pending applications
- Institutional constraints and norms for negotiation

Outputs:

- Leverage assessment: Low, Moderate, or High
- Identification of negotiable versus non-negotiable items
- Success probabilities for each specific negotiation request
- Optimal strategy with timing recommendations
- Expected value improvement from negotiation

Acceptance Criteria:

- System models both parties' incentives and constraints
- System estimates user's replaceability in candidate pool
- System assesses institution's urgency to fill position
- System identifies which items are flexible versus rigid
- System provides bundling recommendations for negotiating multiple items together
- System suggests optimal timing for negotiation initiation
- System quantifies expected value in concrete terms

Derived From: OBJ-1.9, A-2.2

Traces To: Agent capability for game-theoretic negotiation modeling

5.5 Calibration and Self-Knowledge Capabilities

5.5.1 Preference Discovery

Requirement ID: FR-1.5.1

Priority: Medium

Description: The system shall treat user's preferences as unknown variables requiring empirical discovery through exploration.

Inputs:

- User's stated preferences with uncertainty acknowledgment
- Past choices and their outcomes
- Hypothetical scenario responses

Outputs:

- Exploratory action recommendations for testing preferences

- Expected information gain from each exploration
- Preference stability assessment over time
- Recommendations for low-cost preference tests

Acceptance Criteria:

- System identifies which preferences are uncertain
- System suggests experiences that would update self-knowledge
- System balances exploration (learning preferences) versus exploitation (optimizing for known preferences)
- System accounts for changing preferences over lifecycle
- System distinguishes temporary versus stable preference uncertainty

Derived From: OBJ-2.8, A-1.6

Traces To: Agent capability for multi-armed bandit formulation of preference learning

5.5.2 Self-Assessment Calibration

Requirement ID: FR-1.5.2

Priority: Medium

Description: The system shall determine whether user's self-assessment is too harsh (impostor syndrome) or too generous (Dunning-Kruger effect) relative to market reality.

Inputs:

- User's self-rankings across multiple dimensions
- Actual outcome data from application results
- Comparison cohorts with similar profiles

Outputs:

- Calibration correction: over-estimating, accurate, or under-estimating
- Domain-specific assessments identifying which dimensions are miscalibrated
- Confidence in calibration assessment
- Recommended adjustments to self-perception

Acceptance Criteria:

- System compares self-perception to market feedback
- System detects systematic over-estimation or under-estimation
- System controls for selection effects in outcome data

- System identifies domains where metacognition is accurate versus distorted
- System updates calibration as more outcome data arrives
- System distinguishes global versus domain-specific miscalibration

Derived From: OBJ-2.6, A-1.10

Traces To: Agent capability for calibration curve analysis

5.5.3 Strategy Effectiveness Tracking

Requirement ID: FR-1.5.3

Priority: Low

Description: The system shall compare outcomes from different application strategies user has tried to identify personal comparative advantages.

Inputs:

- Application history with strategy labels
- Outcome data for each application
- Opportunity quality controls to isolate strategy effects

Outputs:

- Strategy effectiveness rankings
- Personal comparative advantages in specific approaches
- Strategy-specific success rates with confidence intervals
- Recommendations for strategy adjustments

Acceptance Criteria:

- System tracks outcomes by strategy type
- System controls for opportunity quality differences
- System detects when variance is high (strategy does not matter) versus low (strategy matters)
- System identifies what works for this user specifically not generic advice
- System avoids superstitious causation through proper controls
- System requires minimum sample size before making recommendations

Derived From: OBJ-2.4, A-3.10

Traces To: Agent capability for A-B testing and outcome attribution

5.6 Data Management Requirements

5.6.1 User Data Collection

Requirement ID: FR-1.6.1

Priority: High

Description: The system shall collect and maintain user profile, application history, and longitudinal outcomes with user consent.

Inputs:

- User-provided profile data including credentials, skills, and constraints
- Application submissions and their outcomes over time
- Career milestone events with timestamps
- Decision choices and rationales

Outputs:

- Structured user profile database
- Application history database with outcomes
- Longitudinal outcome records over years
- State change logs tracking profile evolution

Acceptance Criteria:

- System collects data with explicit user consent
- System anonymizes data for aggregate analysis
- System enables longitudinal tracking over multiple years
- System supports counterfactual data structure including choice A versus choice B
- System allows user data export in standard formats
- System allows user data deletion on request
- System maintains data integrity and version control

Derived From: OBJ-2.1, A-1.11

Traces To: Agent capability for state tracking and episodic memory

5.6.2 Public Data Aggregation

Requirement ID: FR-1.6.2

Priority: High

Description: The system shall systematically collect and curate public data sources for institutional and field analysis.

Inputs:

- Institutional websites containing hiring, enrollment, and publication data
- Funding agency reports and strategic plans
- Academic databases including citation and authorship information
- Policy documents and budget proposals

Outputs:

- Institutional performance databases with historical data
- Field vitality metrics updated periodically
- Funding trend databases with temporal granularity
- Advisor and manager performance indicators from public signals

Acceptance Criteria:

- System uses only publicly available data within legal bounds
- System updates on defined schedules: yearly, quarterly, or monthly depending on data type
- System maintains historical time-series not just snapshots
- System validates data quality and detects anomalies
- System documents data provenance for auditability
- System complies with web scraping terms of service

Derived From: OBJ-1.7, A-3.6

Traces To: Agent capability for pre-computed regime analysis

5.6.3 Outcome Data Contribution

Requirement ID: FR-1.6.3

Priority: Medium

Description: The system shall enable users to contribute outcome data years after initial decisions for system learning and network effects.

Inputs:

- User identification through authenticated session
- Outcome descriptions including what happened
- Attribution to specific past decisions
- Intervening factors that affected outcomes

Outputs:

- Enriched outcome database for causal inference

- Improved prediction models through learning
- Network effects data benefiting all users

Acceptance Criteria:

- System provides low-friction outcome reporting interfaces
- System implements privacy-preserving aggregation
- System maintains voluntary participation without coercion
- System enables causal inference dataset growth over time
- System provides value back to contributing users through improved predictions
- System validates outcome data for consistency

Derived From: OBJ-2.15, A-1.12

Traces To: Agent capability for outcome integration and learning

6 Non-Functional Requirements

6.1 Performance Requirements

6.1.1 Response Time

Requirement ID: NFR-2.1.1

Priority: High

Description: The system shall respond to user queries within acceptable latency bounds to maintain engagement.

Specifications:

- Tier 0 operations (staleness checks, filtering): less than 1 second
- Tier 1 operations (cached analysis): less than 5 seconds
- Tier 2 operations (novel analysis): less than 15 seconds
- Tier 3 operations (complex synthesis): less than 30 seconds

Rationale: User engagement requires responsive interaction. Computational efficiency enables fast response despite limited resources.

Acceptance Criteria:

- 95th percentile of Tier 0 operations complete within 1 second
- 95th percentile of Tier 1 operations complete within 5 seconds
- 95th percentile of Tier 2 operations complete within 15 seconds
- 95th percentile of Tier 3 operations complete within 30 seconds
- System provides progress indicators for operations exceeding 5 seconds

Derived From: OBJ-1.11, A-3.5

Traces To: Agent architecture with tiered processing

6.1.2 Throughput

Requirement ID: NFR-2.1.2

Priority: Medium

Description: The system shall support multiple concurrent users within API rate limits.

Specifications:

- Minimum viable: 100 active users per month
- Target capacity: 1,000 active users per month
- Stretch goal: 10,000 active users per month

Rationale: Scalability required for network effects. Must operate within free-tier constraints.

Acceptance Criteria:

- System supports minimum 100 active users with acceptable performance
- System degrades gracefully as user count approaches capacity limits
- System queues requests when rate limits approached
- System provides capacity status to administrators

Derived From: OBJ-2.15, CC-3.1.1

Traces To: Agent coordination with rate limit management

6.2 Resource Constraints

6.2.1 API Token Budget

Requirement ID: NFR-2.2.1

Priority: Critical

Description: The system shall operate within free-tier API token limits as a fundamental design constraint.

Specifications for Primary Model (Groq llama-3.3-70b-versatile):

- Requests per minute (RPM): 30
- Requests per day (RPD): 1,000
- Tokens per minute (TPM): 12,000
- Tokens per day (TPD): 100,000

Monthly Token Allocation Strategy:

- Total monthly budget: approximately 3,000,000 tokens (30 days at 100,000 per day)
- Regime updates (batch processing): 25% (750,000 tokens)
- User queries (reactive): 50% (1,500,000 tokens)
- Learning and calibration: 25% (750,000 tokens)

Per-Query Budget:

- Average query budget: 250 tokens
- Maximum query budget: 1,000 tokens for Tier 3 operations
- Expected queries per month: approximately 6,000

Fallback Models for Specific Tasks:

- meta-llama/llama-4-scout-17b-16e-instruct: Higher TPM (30,000) for batch processing
- llama-3.1-8b-instant: Higher RPD (14,400) for high-frequency filtering

Rationale: Free-tier operation is fundamental design constraint determining all architectural decisions.

Acceptance Criteria:

- Monthly token usage remains below 3,000,000 tokens in normal operation
- Daily token usage remains below 100,000 tokens
- Per-minute token usage remains below 12,000 tokens
- System tracks usage across rolling windows
- System implements automatic throttling when approaching limits
- System maintains 10% buffer for usage spikes

Derived From: A-3.2, OBJ-1.15

Traces To: All agent computational optimization strategies

6.2.2 Hardware Constraints

Requirement ID: NFR-2.2.2

Priority: High

Description: The system shall run on consumer hardware without specialized acceleration.

Specifications:

- CPU: Standard x86_64 or ARM processors
- RAM: 8-16GB maximum usage
- Storage: Less than 10GB for models and data
- No GPU required
- No cloud compute required beyond API calls

Rationale: Accessibility and cost constraints. System must be deployable without infrastructure investment.

Acceptance Criteria:

- System runs on laptop with 8GB RAM
- System does not require GPU acceleration for any operation
- System storage footprint remains below 10GB
- System operates without cloud compute infrastructure
- System performance acceptable on five-year-old hardware

Derived From: A-3.1

Traces To: Agent architecture using lightweight algorithms

6.2.3 Network Constraints

Requirement ID: NFR-2.2.3

Priority: Medium

Description: The system shall minimize network bandwidth usage.

Specifications:

- Batch API calls when possible with minimum five-minute cooldown between batches
- Compress data in transit
- Cache responses locally
- Minimize redundant requests through staleness checking

Rationale: Reduces latency, respects rate limits, enables offline degradation.

Acceptance Criteria:

- API requests batched with minimum five-minute intervals
- Cached data served for repeated queries
- Network usage monitored and logged
- System functions with intermittent connectivity

Derived From: A-3.4

Traces To: Agent message bus with batching coordinator

6.3 Reliability Requirements

6.3.1 Availability

Requirement ID: NFR-2.3.1

Priority: Medium

Description: The system shall remain functional under API service degradation.

Specifications:

- Graceful degradation when API unavailable
- Cached responses as fallback
- Queue non-urgent requests for later processing
- Clear user communication about reduced capability

Acceptance Criteria:

- Core read operations work offline using cached data
- Write operations queued and processed when API available

- User never experiences hard failure always receiving some response
- System communicates degraded mode status clearly
- System automatically recovers when API service resumes

Derived From: OBJ-1.15

Traces To: Agent failure handling and cache-only modes

6.3.2 Data Integrity

Requirement ID: NFR-2.3.2

Priority: High

Description: The system shall maintain accuracy and consistency of stored data.

Specifications:

- User data backed up regularly
- Outcome data validated before incorporation
- Prediction logs maintained for auditing
- Data provenance documented

Acceptance Criteria:

- No data loss on system failure
- Corrupt data detected and flagged
- Historical data maintained as append-only logs
- Data provenance traceable for all analyses
- Automated validation detects inconsistencies

Derived From: OBJ-2.1

Traces To: Agent episodic memory and outcome tracking

6.4 Security and Privacy Requirements

6.4.1 User Privacy

Requirement ID: NFR-2.4.1

Priority: Critical

Description: The system shall protect user privacy and enable data control.

Specifications:

- User data anonymized in aggregate analysis

- No sharing with third parties without explicit consent
- User can export all personal data
- User can delete all personal data
- Differential privacy for aggregate statistics

Compliance Requirements:

- GDPR compliance including data portability and right to deletion
- No sale of user data under any circumstances
- Transparent data usage policies

Acceptance Criteria:

- Aggregated data cannot identify individuals
- User data export functionality tested and verified
- User data deletion removes all traces
- Privacy policy clearly communicated
- Third-party data sharing requires explicit opt-in

Derived From: OBJ-2.14

Traces To: Data management systems with anonymization

6.4.2 Legal Risk Mitigation

Requirement ID: NFR-2.4.2

Priority: High

Description: The system shall avoid legal liability in sensitive assessments.

Specifications:

- Advisor and manager assessments use public signals only
- No defamatory statements
- Confidence intervals on all predictions
- Clear disclaimers that system does not provide legal or financial advice
- Fact-based analysis only without speculation

Acceptance Criteria:

- Legal review of advisor assessment methodology completed
- All claims attributable to observable public data

- No personally identifying information in shared analyses
- Disclaimers presented at appropriate points
- System maintains evidence for all factual claims

Derived From: FR-1.2.3

Traces To: Agent capability for public signal analysis only

6.5 Usability Requirements

6.5.1 Clarity and Interpretability

Requirement ID: NFR-2.5.1

Priority: High

Description: System outputs shall be comprehensible and actionable for target users.

Specifications:

- Clear presentation of uncertainty through confidence intervals
- Transparent reasoning showing logic not just conclusions
- Avoid false precision such as spurious decimal places
- Concrete recommendations providing actionable next steps
- Distinguish facts from inferences

Acceptance Criteria:

- User comprehension testing shows greater than 80% can explain system reasoning
- Actionability testing shows greater than 80% can execute recommendations
- Confidence intervals always accompany probabilistic claims
- Reasoning chains are inspectable by users
- No claims presented without supporting evidence or confidence assessment

Derived From: A-1.1, OBJ-1.4

Traces To: Agent output formatting and explanation generation

6.5.2 Honesty and Calibration

Requirement ID: NFR-2.5.2

Priority: Critical

Description: The system shall provide honest assessments even when uncomfortable.

Specifications:

- No false reassurance or unrealistic optimism
- Explicit acknowledgment of uncertainty
- Uncomfortable truths delivered constructively
- No institutional loyalty bias
- Regular calibration checks comparing predictions versus outcomes

Acceptance Criteria:

- Prediction calibration measured quarterly
- User feedback on honesty versus helpfulness balance collected
- No systematic over-optimism detected in calibration analysis
- System willing to recommend against searching when appropriate
- System willing to critique user's chosen field or institution

Derived From: OBJ-2.12, A-4.4

Traces To: Agent reality feedback and calibration mechanisms

6.6 Maintainability Requirements

6.6.1 Data Update Automation

Requirement ID: NFR-2.6.1

Priority: Medium

Description: The system shall minimize manual curation through automation.

Specifications:

- Automated scraping of public data sources
- Scheduled batch updates: yearly, quarterly, or monthly
- Anomaly detection for data quality issues
- Manual intervention only for exceptions

Acceptance Criteria:

- 90% of data updates fully automated
- Manual review only for flagged anomalies
- Update schedules reliably executed
- Data quality metrics automatically monitored
- Failed updates trigger alerts

Derived From: A-3.14

Traces To: Batch processing schedules for regime updates

6.6.2 Model Update Deployment

Requirement ID: NFR-2.6.2

Priority: Medium

Description: The system shall enable model updates without service interruption.

Specifications:

- A-B testing for model improvements
- Gradual rollout of changes
- Rollback capability if issues detected
- Zero-downtime deployments

Acceptance Criteria:

- Model updates deployed weekly without user disruption
- Performance monitoring detects regressions
- Fast rollback capability under one hour if needed
- A-B test results statistically significant before full rollout

Derived From: OBJ-2.1

Traces To: Agent learning and heuristic compression

6.7 Ethical Requirements

6.7.1 Incentive Alignment

Requirement ID: NFR-2.7.1

Priority: Critical

Description: The system shall optimize for user outcomes not engagement metrics.

Specifications:

- No institutional revenue sources to preserve critique ability
- Success metrics tied to user career outcomes not time on platform
- Willingness to tell users to stop using system if optimal
- Transparent business model using subscription not advertising

Acceptance Criteria:

- Revenue model documented and user-transparent
- No A-B testing for engagement maximization

- Explicit rejection of dark patterns
- System provides stop-searching advice when warranted
- No institutional partnerships that constrain truth-telling

Derived From: OBJ-2.12, A-4.6

Traces To: Agent strategic synthesis willing to recommend cessation

6.7.2 Harm Avoidance

Requirement ID: NFR-2.7.2

Priority: High

Description: The system shall avoid advice that could lead to career catastrophes.

Specifications:

- Explicit flagging of high-risk decisions
- Downside scenario analysis provided
- Conservative confidence intervals on uncertain predictions
- Emergency recommendations such as toxic advisor escape routes

Acceptance Criteria:

- High-risk decisions flagged automatically
- Downside scenarios included in major recommendations
- User feedback on harm avoidance effectiveness positive
- System errs on side of caution for irreversible decisions
- System provides clear warnings for high-uncertainty situations

Derived From: OBJ-1.5

Traces To: Agent risk modeling and downside analysis

7 Computational Constraints

7.1 API Rate Limits

7.1.1 Groq Free-Tier Limits

Constraint ID: CC-3.1.1

Description: System must operate within Groq API free-tier rate limits.

Primary Model: llama-3.3-70b-versatile

Metric	Limit
Requests per minute (RPM)	30
Requests per day (RPD)	1,000
Tokens per minute (TPM)	12,000
Tokens per day (TPD)	100,000

Table 1: Primary Model Rate Limits

Fallback Models:

Model	RPM	TPM	TPD
llama-4-scout-17b-16e	30	30,000	500,000
llama-3.1-8b-instant	30	6,000	500,000

Table 2: Fallback Model Rate Limits

Derived From: A-3.2

Impacts: NFR-2.2.1, All functional requirements requiring LLM calls

7.1.2 Token Budget Distribution

Constraint ID: CC-3.1.2

Description: Monthly token budget must be allocated across system functions.

Total Monthly Budget: 3,000,000 tokens (30 days at 100,000 tokens per day)

Derived From: CC-3.1.1, A-3.12

Impacts: System capacity, feature prioritization, agent design

7.1.3 Per-Query Token Constraints

Constraint ID: CC-3.1.3

Description: Individual user queries must operate within token budgets by tier.

Token Budget by Tier:

Monthly Query Capacity:

- Reactive budget: 1,500,000 tokens

Function	Tokens	Percentage
Regime Updates (Batch)	750,000	25%
Field vitality (yearly)	200,000	
Funding regime (quarterly)	200,000	
Market regime (monthly)	200,000	
Institutional trajectory (quarterly)	150,000	
User Queries (Reactive)	1,500,000	50%
Tier 0 operations	0	
Tier 1 operations	300,000	
Tier 2 operations	800,000	
Tier 3 operations	400,000	
Learning & Calibration	750,000	25%
Outcome integration	300,000	
Model calibration	250,000	
Heuristic compression	200,000	
Total	**3,000,000**	**100%**

Table 3: Monthly Token Budget Allocation

Tier	Token Budget	Expected Frequency
Tier 0	0	95%
Tier 1	50-100	3%
Tier 2	200-400	1.5%
Tier 3	500-1,000	0.5%
Average	**250**	**100%**

Table 4: Per-Query Token Budgets

- Average query cost: 250 tokens
- Monthly capacity: 6,000 queries
- Daily capacity: 200 queries
- Hourly capacity: approximately 8 queries (24/7 operation)

Derived From: CC-3.1.2

Impacts: User capacity, agent tier classification, caching strategy

7.2 Hardware Constraints

7.2.1 Memory Constraints

Constraint ID: CC-3.2.1

Description: System must operate within consumer-grade RAM limits.

Specifications:

- Maximum RAM usage: 8-16GB

- No GPU memory requirements
- Efficient data structures required
- Memory footprint monitoring

Derived From: A-3.1

Impacts: Agent memory architecture, caching strategies, data structure design

7.2.2 Storage Constraints

Constraint ID: CC-3.2.2

Description: System storage footprint must remain minimal.

Specifications:

- Total storage: Less than 10GB
- Pre-computed analyses: Compressed storage
- User data: Efficient serialization
- Historical logs: Periodic compression

Derived From: A-3.1

Impacts: Data retention policies, compression strategies, episodic memory design

7.2.3 Processing Constraints

Constraint ID: CC-3.2.3

Description: System must use CPU-only algorithms without GPU acceleration.

Specifications:

- CPU-only processing for all operations
- Lightweight statistical models preferred
- Algorithmic optimization over computational power
- No deep learning models requiring GPU

Derived From: A-3.1, A-3.11

Impacts: Algorithm selection, model complexity, statistical methods

7.3 Operational Constraints

7.3.1 Batch Processing Requirements

Constraint ID: CC-3.3.1

Description: Regime-level analyses must be pre-computed in batch mode.

Analysis Type	Frequency	Token Budget
Field vitality snapshots	Yearly (September)	200,000
Funding regime analysis	Quarterly	200,000
Market regime detection	Monthly	200,000
Institutional trajectories	Quarterly	150,000
Outcome integration	Weekly	75,000
Model calibration	Monthly	62,500
Heuristic compression	Quarterly	50,000

Table 5: Batch Processing Schedule

Batch Processing Schedule:

Derived From: CC-3.1.2, A-3.4

Impacts: Data staleness, update schedules, caching architecture

7.3.2 Caching Requirements

Constraint ID: CC-3.3.2

Description: Aggressive caching required to minimize API calls.

Caching Strategy:

- State snapshots cached until material changes detected
- Strategic templates cached and reused for 80% of decisions
- Regime analyses cached until scheduled update
- Episodic memories compressed after 30 days at 10:1 ratio
- Semantic memory pruned quarterly removing least-accessed rules

Staleness Tolerance:

- Field vitality: 12 months acceptable staleness
- Funding regimes: 3 months acceptable staleness
- Market regimes: 1 month acceptable staleness
- User state: 14 days maximum staleness

Derived From: CC-3.1.1, A-3.5

Impacts: Data freshness, user experience, agent staleness detection

7.3.3 Request Batching Requirements

Constraint ID: CC-3.3.3

Description: API requests must be batched to respect rate limits.

Batching Strategy:

- Minimum cooldown between batches: 5 minutes
- Accumulate requests until batch threshold or urgent request
- Priority-based request ordering within batches
- Token budget enforcement per batch

Derived From: CC-3.1.1

Impacts: Response latency, user experience, agent communication

8 System Constraints

8.1 Technical Constraints

8.1.1 API Dependencies

Constraint ID: SC-4.1.1

Description: System depends on Groq API availability and free-tier limits.

Implications:

- No fallback to paid tiers (must remain free-tier viable)
- Token budgets are hard constraints not targets
- API unavailability requires graceful degradation
- No guarantee of service availability

Derived From: A-3.2

Impacts: All functional and non-functional requirements

8.1.2 Data Access Constraints

Constraint ID: SC-4.1.2

Description: System limited to public data only.

Implications:

- No proprietary databases
- Web scraping subject to rate limits and terms of service
- Some desirable data sources inaccessible due to paywalls or legal restrictions
- Causal inference dependent on user-contributed outcome data

Derived From: A-3.6, NFR-2.4.2

Impacts: Data quality, feature completeness, prediction accuracy

8.1.3 Computational Constraints

Constraint ID: SC-4.1.3

Description: System limited to algorithms efficient on consumer CPUs.

Implications:

- No GPU acceleration available
- Limited to algorithms that run efficiently on CPU
- Memory footprint bounded by consumer hardware

- Processing speed limited by single-threaded or lightly-threaded algorithms

Derived From: A-3.1

Impacts: Algorithm selection, model complexity, response times

8.2 Business Constraints

8.2.1 Revenue Model Constraints

Constraint ID: SC-4.2.1

Description: System must avoid institutional revenue to preserve independence.

Implications:

- No institutional customers to preserve critique ability
- Subscription-based pricing only no advertising
- Pricing must support development costs while remaining accessible
- No partnerships that constrain truth-telling

Derived From: A-4.6, NFR-2.7.1

Impacts: Business model, feature prioritization, competitive positioning

8.2.2 Legal Constraints

Constraint ID: SC-4.2.2

Description: System must avoid legal liability in assessments.

Implications:

- Must avoid defamation in advisor assessments
- Cannot provide legal or financial advice
- GDPR and privacy regulation compliance required
- Fact-based analysis only without speculation

Derived From: NFR-2.4.2

Impacts: Feature design, output formatting, disclaimer requirements

8.3 User Constraints

8.3.1 User Sophistication Requirements

Constraint ID: SC-4.3.1

Description: System requires sophisticated users capable of probabilistic reasoning.

Implications:

- Target users must be capable of probabilistic reasoning
- System not designed for casual job seekers
- Requires user investment in profile setup and outcome reporting
- May exclude users uncomfortable with uncertainty

Derived From: A-1.1, A-1.3

Impacts: User interface design, explanation complexity, market size

8.3.2 Data Contribution Dependencies

Constraint ID: SC-4.3.2

Description: Network effects depend on voluntary user data contribution.

Implications:

- Network effects depend on voluntary user data contribution
- Outcome data has long lag time measured in years
- Critical mass required for counterfactual analysis
- Early adopters receive less value due to limited data

Derived From: A-1.11, A-3.7

Impacts: Growth strategy, early value proposition, data accumulation timeline

9 Acceptance Criteria

9.1 Functional Acceptance

9.1.1 Minimum Viable Product

MVP Requirements:

The following functional requirements must be operational for MVP launch:

- FR-1.1.2: Search Termination Optimization (functional with Bayesian updating)
- FR-1.2.1: Field Collapse Detection (functional with yearly updates)
- FR-1.2.3: Advisor Quality Prediction (functional with public signal analysis)
- FR-1.3.3: Market Regime Detection (functional with monthly updates)
- FR-1.4.1: Career Path Simulation (functional with basic causal inference)

MVP Acceptance Criteria:

- All MVP functional requirements operational
- Token budget compliance: monthly usage less than 3,000,000 tokens
- Response time: 95th percentile less than 30 seconds
- User comprehension: greater than 70% in testing
- No critical defects

9.1.2 Full Feature Set

Full System Requirements:

- All functional requirements from Section 5 operational
- All non-functional requirements from Section 6 met
- All computational constraints from Section 7 respected
- User acceptance testing passed with target users

Full System Acceptance Criteria:

- 100% of functional requirements operational
- Token budget compliance: 100% of time
- Performance requirements: 95th percentile compliance
- User satisfaction: greater than 75% positive feedback
- No high-priority defects

9.2 Performance Acceptance

9.2.1 Token Budget Compliance

Acceptance Criteria:

- Monthly API usage less than 3,000,000 tokens in 100% of months
- Average query less than 250 tokens in 80% of queries
- Regime updates complete within batch budget 100% of time
- Daily usage less than 100,000 tokens in 95% of days
- Per-minute usage less than 12,000 tokens in 99% of minutes

9.2.2 Response Time Compliance

Acceptance Criteria:

- 95th percentile response time less than 30 seconds
- 50th percentile response time less than 5 seconds
- Tier 0 operations: 99th percentile less than 1 second
- Tier 1 operations: 95th percentile less than 5 seconds
- Tier 2 operations: 95th percentile less than 15 seconds
- Tier 3 operations: 95th percentile less than 30 seconds

9.3 Quality Acceptance

9.3.1 Prediction Calibration

Acceptance Criteria:

- Confidence intervals correctly calibrated (checked quarterly)
- Systematic bias less than 5% in any direction
- User-reported accuracy greater than 70% for verifiable predictions
- Calibration curves updated and validated quarterly
- Prediction logs maintained for all claims

9.3.2 User Satisfaction

Acceptance Criteria:

- User comprehension greater than 80% (can explain system reasoning)
- User actionability greater than 80% (can execute recommendations)

- User trust greater than 70% (willing to follow uncomfortable advice)
- User retention: greater than 60% month-over-month
- User referrals: greater than 30% of new users from existing user referrals

10 Requirements Traceability Matrix

10.1 Assumptions to Objectives

Assumption	Derived Objectives
A-1.6 Users experience preference uncertainty	OBJ-2.8 Detect preference uncertainty and recommend exploration
A-1.10 Users may have poor calibration	OBJ-2.6 Disambiguate impostor syndrome from genuine gaps OBJ-2.10 Correct miscalibration in self-assessment
A-1.11 Users will contribute outcome data	OBJ-2.1 Improve predictions through outcome feedback OBJ-2.15 Create network effects through contributions
A-2.1 Opportunities arrive sequentially	OBJ-1.2 Minimize wasted effort on low-probability opportunities
A-2.3 Prestige carries excess weight	OBJ-2.11 Provide capabilities competitors cannot
A-2.6 Markets shift between regimes	OBJ-1.6 Detect market regime shifts early
A-2.7 Fields follow trajectories	OBJ-1.7 Identify field and institutional trajectories
A-2.8 Public signals correlate with quality	OBJ-1.8 Predict advisor quality from public signals
A-3.2 Free-tier API limits	OBJ-1.11 Minimize token consumption OBJ-1.12 Avoid unnecessary recursion OBJ-1.15 Operate within free-tier constraints
A-3.4 Pre-computation viable	OBJ-1.14 Maximize use of pre-computed analyses
A-3.9 Causal inference possible	OBJ-1.10 Model counterfactual career paths
A-4.4 Computational honesty differentiates	OBJ-2.12 Maintain incentive alignment
A-4.6 Competitors have misalignments	OBJ-2.11 Provide structurally impossible capabilities OBJ-2.14 Resist institutional capture

10.2 Objectives to Requirements

Objective	Implementing Requirements
OBJ-1.2 Minimize wasted effort	FR-1.1.2 Search Termination Optimization FR-1.3.3 Market Regime Detection

Objective	**Implementing Requirements**
OBJ-1.3 Preserve option value	FR-1.1.1 Path Irreversibility Analysis FR-1.1.3 Decision Sequencing Optimization
OBJ-1.4 Enable informed decisions	FR-1.4.2 Information Gap Prioritization NFR-2.5.1 Clarity and Interpretability
OBJ-1.5 Reduce career mistakes	FR-1.2.1 Field Collapse Detection FR-1.2.3 Advisor Quality Prediction NFR-2.7.2 Harm Avoidance
OBJ-1.6 Detect regime shifts early	FR-1.2.1 Field Collapse Detection FR-1.2.2 Funding Regime Shift Prediction FR-1.3.3 Market Regime Detection
OBJ-1.7 Identify trajectories	FR-1.3.2 Institutional Trajectory Forecasting
OBJ-1.8 Predict advisor quality	FR-1.2.3 Advisor and Manager Quality Prediction
OBJ-1.9 Calculate trade-offs	FR-1.4.3 Negotiation Leverage Analysis
OBJ-1.10 Model counterfactuals	FR-1.4.1 Career Path Simulation
OBJ-1.11 Minimize tokens	NFR-2.2.1 API Token Budget CC-3.1.1 Groq Free-Tier Limits
OBJ-1.15 Operate within constraints	NFR-2.2.1 API Token Budget NFR-2.3.1 Availability
OBJ-2.1 Improve through feedback	FR-1.6.3 Outcome Data Contribution NFR-2.6.2 Model Update Deployment
OBJ-2.6 Disambiguate impostor syndrome	FR-1.5.2 Self-Assessment Calibration
OBJ-2.8 Detect preference uncertainty	FR-1.5.1 Preference Discovery
OBJ-2.11 Provide unique capabilities	FR-1.3.1 Prestige Arbitrage Identification
OBJ-2.12 Maintain alignment	NFR-2.7.1 Incentive Alignment NFR-2.5.2 Honesty and Calibration
OBJ-2.14 Resist capture	NFR-2.4.1 User Privacy SC-4.2.1 Revenue Model Constraints
OBJ-2.15 Create network effects	FR-1.6.3 Outcome Data Contribution

10.3 Requirements to Agent Capabilities

Requirement	Agent Capabilities
FR-1.1.1 Path Irreversibility	State Estimator: Track state transitions Strategy Synthesizer: Model future opportunity sets Temporal Coordinator: Calculate reversal timelines
FR-1.1.2 Search Termination	Strategy Synthesizer: Bayesian stopping rules Reality Feedback: Outcome tracking for calibration State Estimator: Application history tracking
FR-1.1.3 Decision Sequencing	Temporal Coordinator: Deadline topology analysis Temporal Coordinator: Constraint satisfaction scheduling Information Gap: Uncertainty resolution timing
FR-1.2.1 Field Collapse	Pre-computed regime analysis (batch) Longitudinal trend detection Leading indicator monitoring
FR-1.2.2 Funding Regime Shifts	Pre-computed policy analysis (batch) Strategic plan interpretation Forward-looking intelligence extraction
FR-1.2.3 Advisor Quality	Public signal pattern detection Statistical anomaly identification Diagnostic question generation (Information Gap)
FR-1.3.1 Prestige Arbitrage	Outcome metrics comparison Prestige-quality gap detection Strategy Synthesizer: Arbitrage recommendations
FR-1.3.2 Institutional Trajectory	Time-series analysis (batch) Trajectory classification Risk factor identification
FR-1.3.3 Market Regime	Aggregate indicator monitoring (batch) Regime change detection Strategy Synthesizer: Regime-appropriate tactics
FR-1.4.1 Career Path Simulation	Causal inference with synthetic controls Outcome database matching Strategy Synthesizer: Persist-pivot recommendations
FR-1.4.2 Information Gap	Information Gap: Uncertainty enumeration Information Gap: EVOI calculation Information Gap: Diagnostic question design
FR-1.4.3 Negotiation Leverage	Strategy Synthesizer: Game-theoretic modeling State Estimator: BATNA assessment Opportunity Filter: Replaceability estimation

Requirement	Agent Capabilities
FR-1.5.1 Preference Discovery	State Estimator: Preference tracking Strategy Synthesizer: Exploration-exploitation balance Reality Feedback: Preference stability assessment
FR-1.5.2 Self-Assessment	Reality Feedback: Calibration analysis State Estimator: Self-perception tracking Outcome comparison to market feedback
FR-1.5.3 Strategy Effectiveness	Reality Feedback: Outcome attribution Reality Feedback: Strategy comparison Heuristic learning from successful patterns
FR-1.6.1 User Data Collection	State Estimator: Profile management State Estimator: Change log maintenance Privacy-preserving storage
FR-1.6.2 Public Data Aggregation	Batch data collection scripts Systematic curation pipelines Data quality validation
FR-1.6.3 Outcome Contribution	Reality Feedback: Outcome integration Network effects data accumulation Causal inference dataset growth
NFR-2.1.1 Response Time	Tiered processing architecture Aggressive caching (State Estimator) Pre-computation (batch analyses)
NFR-2.2.1 Token Budget	All agents: Token-aware design Tier 0 operations: Zero-token algorithms Message bus: Batching coordination
NFR-2.3.1 Availability	State Estimator: Cache-only mode Strategy Synthesizer: Degraded synthesis Request queueing system
NFR-2.5.2 Honesty	Reality Feedback: Calibration enforcement Strategy Synthesizer: Truth-telling mandate No engagement optimization
NFR-2.7.1 Incentive Alignment	Strategy Synthesizer: Outcome optimization No institutional revenue constraints Willingness to recommend cessation

11 Glossary

BATNA Best Alternative To Negotiated Agreement; the user's fallback option if current negotiation fails

Bayesian Updating
Statistical method for updating probability estimates as new evidence arrives

Causal Inference
Statistical methods for determining cause-and-effect relationships rather than mere correlations

Counterfactual
Hypothetical scenario describing what would have happened under different choices

Episodic Memory
System memory of specific events and outcomes with temporal context

EVOI Expected Value Of Information; the value gained by resolving uncertainty before making a decision

Goodhart's Law
When a measure becomes a target, it ceases to be a good measure; metrics degrade when optimized

Hysteresis Property where switching threshold exceeds zero, creating resistance to reversing decisions

Option Value The value of preserving future flexibility and maintaining multiple accessible paths

Pooling Equilibrium
Market state where high-quality and low-quality participants become indistinguishable

Prestige Arbitrage
Exploiting mismatch between institutional prestige and actual opportunity quality

Regime (Market)
Prevailing labor market conditions: buyer's market (employer advantage) versus seller's market (candidate advantage)

Reservation Utility
Minimum acceptable outcome for user; below this threshold, user continues searching

RPD Requests Per Day; API rate limit on daily request volume

RPM Requests Per Minute; API rate limit on per-minute request frequency

Satisficing Accepting "good enough" rather than searching for optimal; bounded rationality strategy

Secretary Problem
Optimization problem of selecting best option from sequential arrivals with no recall

Semantic Memory
: System memory of general knowledge, patterns, and rules without temporal context

Signaling
: Costly investment to credibly communicate private information about quality

Staleness
: Degree to which cached data diverges from current reality

Synthetic Control
: Comparison group constructed through matching to estimate counterfactual outcomes

TPD
: Tokens Per Day; API rate limit on daily token consumption

TPM
: Tokens Per Minute; API rate limit on per-minute token consumption

Type I Error
: False positive; accepting bad candidate (institution's primary concern)

Type II Error
: False negative; rejecting good candidate (candidate bears asymmetric cost)

Winner's Curse
: Phenomenon where auction winners often overpay; those most eager may most misestimate value

Working Memory
: System memory of current active state and ongoing processes

12 References

12.1 Problem Domain Literature

1. Spence, M. (1973). Job Market Signaling. *Quarterly Journal of Economics*, 87(3), 355-374.

2. Akerlof, G. A. (1970). The Market for Lemons: Quality Uncertainty and the Market Mechanism. *Quarterly Journal of Economics*, 84(3), 488-500.

3. Simon, H. A. (1956). Rational Choice and the Structure of the Environment. *Psychological Review*, 63(2), 129-138.

4. Ferguson, T. S. (1989). Who Solved the Secretary Problem? *Statistical Science*, 4(3), 282-289.

5. Kahneman, D., & Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, 47(2), 263-291.

12.2 Technical Methods Literature

1. Pearl, J. (2009). *Causality: Models, Reasoning, and Inference* (2nd ed.). Cambridge University Press.

2. Rubin, D. B. (2005). Causal Inference Using Potential Outcomes. *Journal of the American Statistical Association*, 100(469), 322-331.

3. Abadie, A., & Gardeazabal, J. (2003). The Economic Costs of Conflict: A Case Study of the Basque Country. *American Economic Review*, 93(1), 113-132.

4. Gelman, A., et al. (2013). *Bayesian Data Analysis* (3rd ed.). CRC Press.

5. Sutton, R. S., & Barto, A. G. (2018). *Reinforcement Learning: An Introduction* (2nd ed.). MIT Press.

12.3 System Architecture Literature

1. Wooldridge, M. (2009). *An Introduction to MultiAgent Systems* (2nd ed.). John Wiley & Sons.

2. Russell, S., & Norvig, P. (2020). *Artificial Intelligence: A Modern Approach* (4th ed.). Pearson.

3. Dwork, C., & Roth, A. (2014). The Algorithmic Foundations of Differential Privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3-4), 211-407.