



**THE CALIFORNIA CASINO SCORE (CAC SCORE):  
A STRATIFIED SURVEY-AND-TESTING FRAMEWORK  
FOR EVALUATING OFFSHORE ONLINE CASINOS  
AMONG CALIFORNIA PLAYERS**

*A Research Study*

Conducted by

**Derek Loomis · Chloe Marsh · Ryan Carrigan  
Joanna Pham · Aaron Whitfield · Theo Ashworth**

Palo Alto Casino Analysis Collective  
Palo Alto, California

**2026**

*N = 4,217 verified California players (21+) · 95% confidence · ±1.5% margin of error*

## Abstract

---

This thesis introduces and validates the California Casino Score (CAC Score), a stratified survey-and-testing framework for evaluating offshore online casinos serving California players. Existing consumer casino ratings rely on opaque, single-number star scores that are rarely grounded in player data, regionally specific, or methodologically transparent. To address this gap, the study fielded a verified survey of  $N = 4,217$  California residents aged 21 and over who had played real-money online casino games within the preceding twelve months (January–April 2026), a sampling fraction of approximately one per cent of the estimated active audience, yielding a 95% confidence level with a  $\pm 1.5\%$  margin of error. Respondents were recruited through a stratified online panel (78%) and managed snowball referral (22%), with regional quotas proportional to California's player distribution and layered verification including government-ID age checks and funded-account screening. A 64-item instrument across eight weighted domains demonstrated strong internal-consistency reliability (Cronbach's  $\alpha = 0.89$ ). The eight domains—anchored in the gambling-trust, payout, and composite-indicator literatures—are combined through a transparent, reproducible weighting formula into a single 0–100 composite. Across fifteen evaluated operators, CAC Scores ranged from 98 (Ignition) to 75 (VoltageBet), with crypto-forward banking emerging as the strongest discriminator of payout performance ( $t(13) = 3.31, p = .006, d = 1.71$ ) and regional variation proving statistically significant but practically negligible ( $\eta^2 = .0035$ ). A key methodological finding is that the conventional five-point star rating compresses this 23-point range onto only five discrete steps, discarding decision-relevant gradation that the multidimensional CAC Score preserves. The framework offers players a transparent alternative to affiliate-driven rankings and demonstrates that a survey-grounded composite index can be maintained at consumer scale without sacrificing methodological defensibility.

**Keywords:** *online gambling; casino ratings; California; survey methodology; stratified sampling; composite indicator; Cronbach's alpha; consumer protection; responsible gambling; CAC Score*

## Table of Contents

---

### Chapter 1 — Introduction

- 1.1 Background
- 1.2 Problem Statement
- 1.3 Research Aim and Objectives
- 1.4 Research Questions and Hypotheses
- 1.5 Significance and Contribution
- 1.6 Scope and Delimitations
- 1.7 Definition of Key Terms
- 1.8 Thesis Structure

### Chapter 2 — Literature Review

- 2.1 Introduction to the review
- 2.2 Online gambling research and consumer behaviour
- 2.3 Existing casino rating and review methodologies and their limitations
- 2.4 Survey methodology and sampling theory
- 2.5 Measurement: Likert scales, instrument design, validity and reliability
- 2.6 Trust, licensing, payout and player-protection literature
- 2.7 California regulatory and market context
- 2.8 Identified gaps in the literature
- 2.9 Conceptual framework

### Chapter 3 — Methodology

- 3.1 Research design
- 3.2 Population and sampling frame
- 3.3 Sampling strategy
- 3.4 Recruitment and participant verification
- 3.5 Funded-account screening and the per-casino rule
- 3.6 Instrument design
- 3.7 Reliability and validity
- 3.8 The CAC Score weighting model and formula
- 3.9 Statistical analysis plan
- 3.10 Ethical considerations

### Chapter 4 — Results and Data

- 4.1 Overview of Fieldwork and Response
- 4.2 Sample Demographics
- 4.3 Descriptive Results by Pillar
- 4.4 Per-Casino Results
- 4.5 Inferential Results
- 4.6 Reliability Results

#### 4.7 Summary of Key Findings

### **Chapter 5 — Discussion**

- 5.1 Interpretation of Principal Findings
- 5.2 The CAC Score Versus the Editorial Star Rating
- 5.3 What California Players Prioritise
- 5.4 Per-Casino Insights
- 5.5 Regional and Demographic Patterns
- 5.6 Implications for Players, Operators and the Rating Industry
- 5.7 Comparison With Prior Literature

### **Chapter 6 — Conclusion, Limitations and Future Research**

- 6.1 Summary of the Study and Responses to the Research Questions
- 6.2 Contributions
- 6.3 Limitations
- 6.4 Recommendations
- 6.5 Future Research
- 6.6 Concluding Remarks

### **References**

### **Appendices**

- Appendix A — Survey Instrument
- Appendix B — Per-Casino Data Appendix
- Appendix C — Verification & Screening Protocol
- Appendix D — Regional Quota Detail & Weighting Worksheet
- Appendix E — Responsible-Gambling Resources Provided to Participants

## List of Tables

---

- Table 1.1 — Definition of key terms used throughout the thesis
- Table 1.2 — The eight CAC Score pillars and their weights (summing to 100)
- Table 2.1 — Information loss: the five-point star rating versus the multidimensional CAC Score (study dataset, sorted by CAC Score)
- Table 2.2 — Comparison of rating methodologies against four limitation criteria
- Table 2.3 — Conventional Cronbach's alpha interpretation bands and the study's reported reliability
- Table 3.1 — Regional strata and proportional sampling quotas
- Table 3.1b — Secondary monitoring quotas (age band, gender, primary device)
- Table 3.2 — The eight pillars: example items, response format, and composite weights
- Table 3.3 — Worked CAC Score calculation for BetOnline
- Table 4.1 — Response and verification funnel (January–April 2026)
- Table 4.2 — Regional distribution of respondents (N = 4,217)
- Table 4.3 — Age-band distribution (N = 4,217)
- Table 4.4 — Gender distribution (N = 4,217)
- Table 4.5 — Primary device for online casino play (N = 4,217)
- Table 4.6 — Pillar weights and aggregate sub-score descriptives across fifteen operators
- Table 4.7 — Master per-casino results, sorted by CAC Score descending
- Table 4.8 — Independent-samples t-test, Payout Speed & Banking sub-score
- Table 4.9 — One-way ANOVA of respondent-level CAC Score across six regions
- Table 4.10 — Internal-consistency reliability (Cronbach's alpha) by domain
- Table 5.1 — Star-rating bands and the CAC-Score ranges they conceal
- Table A1 — Screening and verification block (routed out on failure)
- Table A2 — California Player Survey items (7-point Likert unless noted)
- Table A3 — Trust & Licensing items
- Table A4 — Payout Speed & Banking items
- Table A5 — Bonuses & Value items
- Table A6 — Game Selection items
- Table A7 — Security & Fairness items
- Table A8 — Customer Support items
- Table A9 — Mobile & Responsible Gambling items
- Table B1 — Master per-casino data matrix (all 15 casinos; pillar sub-scores 0–100; sorted by CAC Score, descending)
- Table C1 — Sequential verification and screening funnel
- Table D1 — Regional strata, proportional quotas, and target counts (N = 4,217)
- Table D2 — The eight weighted pillars (weights sum to 100)
- Table D3 — Worked weighting example: Ignition (top-ranked, CAC Score 98)
- Table E1 — Responsible-gambling resources surfaced to all participants

## List of Figures

---

Figure 1.1 — Relative pillar weights in the CAC Score composite, illustrating the dominance of the California Player Survey and Trust & Licensing pillars.

Figure 2.1 — Schematic relationship between sample size and margin of error at 95% confidence, showing diminishing precision returns and the position of the study's achieved sample ( $N = 4,217$ ,  $MoE \approx \pm 1.5\%$ ).

Figure 2.2 — Conceptual framework of the CAC Score: eight literature-anchored pillars (weights summing to 100) feeding a single 0–100 composite, with verified-player gating on the survey pillar and effect-size-reported regional analysis.

Figure 3.1 — Convergent mixed-methods design: parallel survey and structured-testing streams integrated at the pillar-scoring stage to produce a single composite CAC Score per casino.

Figure 4.1 — Response and verification funnel from invitation to final verified sample.

Figure 4.2 — Regional distribution of respondents (six strata).

Figure 4.3 — Age-band distribution of respondents.

Figure 4.4 — Primary device for online casino play.

Figure 4.5 — Mean pillar sub-scores with standard-deviation whiskers across fifteen casinos.

Figure 4.6 — CAC Score versus editorial star rating by casino.

Figure 4.7 — Verified-player counts by casino, with the 100-respondent validity threshold marked.

# Chapter 1 — Introduction

---

## 1.1 Background

Online casino gambling occupies an unusual and largely unregulated position within California, the most populous state in the United States. Unlike states such as New Jersey, Pennsylvania, Michigan, Connecticut, and West Virginia, which have enacted statutory frameworks for licensed online casino gaming (commonly abbreviated as iGaming), California has not authorized any form of real-money online casino play. The state permits tribal and cardroom table games, a state lottery, and pari-mutuel horse wagering, yet it has consistently declined to extend legalization to interactive online casino products. Repeated legislative and ballot efforts, including the high-profile sports-wagering propositions placed before voters in November 2022, were decisively rejected, leaving California without a regulated online gambling market of any kind (Eilers & Krejcik Gaming, 2023). The practical consequence is that the demand which exists in this market, and the demand is substantial, is served almost entirely by offshore operators licensed in jurisdictions such as Curaçao, Panama, and Costa Rica, operating in what is best described as a grey market.

The scale of this grey-market activity is frequently underestimated. Drawing on panel-derived prevalence estimates triangulated against operator traffic data, the Palo Alto Casino Analysis Collective estimates that approximately **420,000** California residents actively played real-money online casino games in the twelve months preceding this study. This audience does not represent a fringe population. It reflects a sustained behavioural pattern among adults who, in the absence of a domestic regulated option, transact with offshore platforms that accept California players. These operators advertise through affiliate networks, accept cryptocurrency and conventional payment instruments, and frequently market themselves using national or generic terminology that obscures the jurisdictional realities a California player actually faces (American Gaming Association, 2023).

California differs from regulated states in ways that materially affect the player experience and, by extension, the validity of any evaluation framework applied to it. In a regulated state, a player can consult a public licensing register, rely on a state-mandated dispute-resolution mechanism, and assume that an operator has posted bonds and submitted to independent testing. None of these protections exist for the California player engaging an offshore operator. Licensing claims cannot be verified against a domestic authority; recourse in the event of a withheld withdrawal is limited; and the responsible-gambling infrastructure that accompanies regulated markets is inconsistently implemented offshore (Gainsbury, 2015). The California player therefore operates in an information environment characterized by

asymmetry, where the operator holds disproportionate knowledge of its own practices and the player must rely on secondary signals to assess trustworthiness.

A further distinguishing feature concerns the legal age of participation. Whereas many international jurisdictions and several discussions of online gambling default to an eighteen-year minimum, the prevailing minimum age for casino-style gambling in California is twenty-one. This study adopts the **21+** framing throughout, both as a methodological screening criterion and as an ethical commitment. All survey respondents were age-verified to twenty-one or older through government-identification verification, and no participant below that threshold was admitted to the sample. This framing aligns the research with the operative norms of California gambling participation and ensures that the resulting evaluations describe an adult population entitled, under the prevailing age convention, to engage in this activity (Williams & Volberg, 2021).

The demographic and behavioural composition of the California online-casino audience is also distinctive. Among the verified respondents in this study, mobile devices were the primary mode of access for sixty-four percent of players, with desktop accounting for twenty-eight percent and tablet for eight percent. This mobile-first orientation has direct implications for how operators should be evaluated, since a platform that performs adequately on desktop but poorly on a small touchscreen serves the majority of California players badly. The audience is geographically concentrated in Southern California, which accounts for fifty-eight percent of the population frame, followed by the San Francisco Bay Area at twenty percent, with the Central Valley, Sacramento Metro, Central Coast, and North State regions comprising the remainder. These regional concentrations matter because payment preferences, connectivity, and even regulatory awareness vary across the state, and a credible evaluation must account for this heterogeneity rather than treating California as a monolith.

### ***1.1.1 The offshore and grey-market reality***

The grey market is not a temporary anomaly that pending legislation will soon resolve. The defeat of the 2022 ballot measures, combined with the entrenched and sometimes competing interests of tribal gaming operators, commercial sportsbooks, and cardrooms, has produced a political stalemate that observers expect to persist for several years (Eilers & Krejcik Gaming, 2023). In the interim, California players continue to fund accounts at offshore casinos, and the volume of that activity continues to grow. The relevant question for player-protection research, therefore, is not whether Californians will gamble online, because they demonstrably already do, but rather how they can be helped to distinguish operators that

treat them fairly from operators that do not. This reframing motivates the entire study and explains why a California-specific evaluation instrument is both necessary and overdue.

Offshore operators serving California typically hold licenses from jurisdictions whose regulatory intensity varies widely. A Curaçao master or sublicense, for example, imposes obligations that differ substantially from those a regulated United States authority would impose, and the enforcement capacity behind such licenses is limited from the perspective of a player located thousands of miles away. This does not render every offshore operator untrustworthy; many maintain sound practices, honour withdrawals promptly, and implement responsible-gambling tools voluntarily. The difficulty for the California player is one of discrimination: in the absence of a domestic register or a state-backed dispute mechanism, the player cannot easily separate the conscientious operator from the negligent one. The signals that would ordinarily perform this sorting function in a regulated market are simply unavailable, and the signals that remain, such as marketing claims and unverified review aggregates, are precisely those most susceptible to manipulation (Auer & Griffiths, 2017).

Cryptocurrency has further reshaped the California offshore experience. A substantial share of offshore deposits and withdrawals now move through digital assets, which alter both the speed and the risk profile of banking. For some players, crypto rails deliver faster payouts than conventional methods can achieve under offshore conditions; for others, they introduce volatility, custody concerns, and an additional layer of technical friction. Any framework that aspires to describe the California experience accurately must treat banking and payout performance as a first-order concern rather than an afterthought, because for the mobile-first, grey-market California player, the question of whether and how quickly winnings can actually be retrieved is among the most consequential a platform must answer.

## **1.2 Problem Statement**

The central problem addressed by this study is that the casino-rating resources available to California players are systematically misaligned with the conditions those players actually face. The dominant rating systems in the English-language market are produced for global or pan-national audiences and are calibrated to the realities of regulated jurisdictions. They privilege licensing badges issued by authorities that have no force in California, they assume dispute-resolution pathways that California players cannot access, and they aggregate player sentiment from populations whose payment methods, connectivity, and legal context differ markedly from those of a Sacramento or San Diego resident. A national five-star score, in short, answers a question the California player did not ask.

This misalignment manifests in several concrete failures. First, generic ratings overweight signals that are legible to the rating provider but operationally meaningless to the California player, such as licensing in a jurisdiction whose register the player can neither read nor enforce. Second, they underweight or omit entirely the factors that most directly determine a California player's lived experience, particularly payout speed under offshore banking conditions, cryptocurrency support, and mobile usability for a mobile-first audience. Third, where these systems incorporate player feedback at all, they rarely verify that the feedback originates from people who actually funded real-money accounts at the operator in question, leaving the resulting scores vulnerable to incentivized reviews, affiliate manipulation, and the sentiment of players in entirely different regulatory environments (Auer & Griffiths, 2017).

The consequence is a credibility gap. A California player consulting a conventional review site may encounter an operator awarded a perfect five-star rating, yet that rating reflects neither the experience of California players specifically nor any verification that reviewers had genuine funded exposure to the platform. The present study finds direct evidence of this disconnect within its own data. Several operators that carry maximal or near-maximal conventional star ratings score materially lower on the California-calibrated framework once California player verification and California-relevant pillar weighting are applied, while at least one operator carrying a middling three-star conventional rating performs strongly on the California-specific measure. This divergence is not noise; it is the predictable result of evaluating California players with instruments designed for someone else. The problem this study sets out to solve is the absence of a rigorous, transparent, and California-specific framework that closes that gap.

The problem is compounded by structural incentives within the affiliate ecosystem that surrounds online casino marketing. Many of the rating resources a California player is most likely to encounter are operated by parties with commercial relationships to the operators they rate, and the resulting scores can drift toward promotion rather than assessment. Even where a rating provider intends to be impartial, the absence of a verification mechanism means that the underlying signal can be contaminated by reviews from non-players, by sentiment imported from regulated markets with entirely different conditions, or by coordinated efforts to inflate or depress a competitor's standing (Gainsbury, Russell, & Hing, 2014). The California player, who cannot readily detect any of this, is left to act on information whose provenance is opaque. A framework that does not verify funded exposure cannot claim to describe player experience; it can only describe the aggregate of whatever sentiment happened to be submitted, weighted by whatever undisclosed scheme the provider chose.

There is, finally, a temporal dimension to the misalignment. Generic rating systems are frequently slow to reflect changes in operator behaviour, and a score earned under one set of banking or bonus conditions may persist long after those conditions have deteriorated. For a population whose experience is dominated by payout reliability and mobile usability, both of which can change rapidly, a stale or non-California score is not merely imprecise but potentially misleading. The problem this study confronts is therefore not a single deficiency but a cluster of interacting failures, jurisdictional, demographic, methodological, commercial, and temporal, all of which point to the same conclusion: California players need an instrument built for them.

### **1.3 Research Aim and Objectives**

The overarching aim of this study is to develop, validate, and apply a California-specific framework for evaluating offshore online casinos as they are experienced by California residents aged twenty-one and older, and to demonstrate that such a framework yields materially different and more defensible assessments than conventional national or generic rating systems. The framework introduced for this purpose is the **California Casino Score (CAC Score)**, a single composite measure on a 0–100 scale derived from eight weighted evaluation pillars.

In service of this aim, the study pursues the following specific objectives:

1. To characterize the population of California residents aged twenty-one and older who played real-money online casino games in the preceding twelve months, including their regional, demographic, device, and behavioural profiles.
2. To design a multi-pillar evaluation instrument that integrates verified California player sentiment with independent expert testing across trust, banking, bonus, game, security, support, and mobile dimensions.
3. To establish and apply a funded-account verification procedure that ensures California player sub-scores rest only on respondents with genuine real-money exposure to a named operator.
4. To compute CAC Scores for a sample of offshore operators serving California players and to assign interpretable performance bands to those scores.
5. To compare CAC Scores against conventional star ratings and quantify the extent of divergence between California-specific and generic evaluations.
6. To assess the internal-consistency reliability and statistical robustness of the instrument and to surface the player-protection implications of the findings.

## 1.4 Research Questions and Hypotheses

The study is organized around five research questions, each accompanied by one or more testable hypotheses where the question is amenable to statistical evaluation.

**RQ1.** What are the defining characteristics of California residents aged twenty-one and older who engage in real-money online casino play, and how do these characteristics differ across the state's regions?

*H1.* Player priorities and platform experiences differ significantly across California's regional strata, such that a one-way analysis of variance across regions will reveal statistically significant differences in at least one key evaluation dimension.

**RQ2.** Can a multi-pillar, weighted composite measure reliably capture the dimensions of operator quality that matter to California players?

*H2.* The eight-pillar instrument will demonstrate acceptable to strong internal-consistency reliability, operationalized as an overall Cronbach's alpha at or above the conventional 0.70 threshold.

**RQ3.** Does evaluation based on verified California player data and California-relevant pillar weighting produce assessments that differ materially from conventional national or generic star ratings?

*H3.* The rank ordering of operators produced by the CAC Score will diverge substantially from the rank ordering implied by conventional star ratings, with multiple operators shifting bands between the two systems.

**RQ4.** Does requiring funded-account verification change which operators qualify for survey-based assessment, and how does this requirement affect the validity of the resulting scores?

*H4.* A non-trivial proportion of operators serving California players will fail to attain the minimum verified-player threshold for survey-based scoring and will consequently be assessed on tester evidence alone, demonstrating the discriminating effect of the verification rule.

**RQ5.** How do offshore operators serving the California market distribute across the CAC Score performance bands, and what does that distribution reveal about the overall quality of options available to California players?

These questions move deliberately from description (RQ1) through instrument construction and validation (RQ2), to the central comparative claim of the study (RQ3), the integrity of the verification

mechanism (RQ4), and finally the substantive market characterization (RQ5). Together they structure the empirical chapters that follow.

## 1.5 Significance and Contribution

The principal contribution of this study is the introduction of the California Casino Score (CAC Score), the first evaluation framework known to the authoring organization that is purpose-built for California players engaging offshore online casinos. The framework's significance rests on three design commitments that distinguish it from prevailing rating systems.

First, the CAC Score privileges **verified California player sentiment** as its single most heavily weighted input. The California Player Survey pillar carries a twenty-percent weight, the largest of any pillar, and contributes only where at least one hundred verified respondents funded a real-money account at the named operator within the preceding twelve months. This requirement directly addresses the verification failure that undermines conventional ratings, ensuring that the loudest signal in the composite originates from Californians with genuine financial exposure to the platform being judged.

Second, the framework is **transparent and reproducible by construction**. The eight pillars, their weights, the 0–100 sub-score scale, the linear aggregation formula, and the interpretive bands are all stated explicitly, so that any analyst applying the same inputs would arrive at the same score. This stands in contrast to proprietary star systems whose internal weighting is rarely disclosed and cannot be independently audited.

Third, the framework is **California-calibrated** rather than globally generic. Its pillar weights reflect what matters to a mobile-first, grey-market, twenty-one-and-older California audience operating without domestic regulatory recourse, allocating substantial weight to payout speed and banking, bonus value, and game selection precisely because these factors govern the offshore California experience.

The practical significance follows directly. By demonstrating that California-specific evaluation produces materially different assessments from conventional ratings, the study equips California players to make better-informed decisions, supplies player-protection advocates and researchers with a replicable instrument, and offers a transferable template that could be adapted to other grey-market jurisdictions facing analogous information asymmetries (Gainsbury, Russell, & Hing, 2014). The study is conducted independently of any operator and surfaces responsible-gambling resources to all participants, reinforcing that its purpose is player protection rather than promotion.

## **1.6 Scope and Delimitations**

The scope of this study is bounded deliberately, and several delimitations follow from the research design. The target population comprises California residents aged twenty-one and older who played real-money online casino games within the twelve months preceding the study; players resident outside California, those under twenty-one, and those who did not play real money are outside scope. Fieldwork was confined to the January through April 2026 window, and the findings should be read as a snapshot of that period rather than a continuous longitudinal account.

The study evaluates a defined set of offshore operators that accept California players; it does not assess regulated domestic operators, because none exist for online casino play in California, nor does it evaluate sweepstakes or social-casino products, which fall under a different legal and behavioural category. The instrument is a self-administered online questionnaire, which delimits the population to players reachable through online panels and referral and able to complete a web-based instrument; players who gamble offshore but are not reachable by these means are not represented.

A defining delimitation concerns the funded-account rule. Survey-based California Player Survey sub-scores are assigned only to operators that attracted at least one hundred verified funded respondents within the relevant window. Operators below that threshold are assessed on tester evidence and explicitly flagged as having limited player data. This rule deliberately constrains the survey component of the framework in exchange for validity, and readers should interpret survey-derived sub-scores only for operators that cleared the threshold. Finally, the sampling fraction of approximately one percent is computed against the estimated active online-casino audience of roughly 420,000, not against the general California population; this distinction is material to any interpretation of the sample's representativeness and is treated as a formal delimitation of the claims the study makes.

## **1.7 Definition of Key Terms**

The following terms carry specific technical meanings throughout this thesis and are defined here to avoid ambiguity in subsequent chapters.

**Table 1.1 — Definition of key terms used throughout the thesis**

Term	Definition as used in this study
<b>CAC Score</b>	The California Casino Score, a composite measure reported on a 0–100 scale, computed as the sum of each pillar's 0–100 sub-score multiplied by that pillar's weight, divided by 100. Interpreted by band: 85 and above Excellent, 75–84 Good, 65–74 Fair, below 65 Weak.
<b>Pillar</b>	One of the eight weighted evaluation dimensions that compose the CAC Score. Each pillar receives a 0–100 sub-score and carries a fixed weight summing across all pillars to 100.
<b>Pillar weight</b>	The fixed percentage contribution a pillar makes to the composite CAC Score, set to reflect the relative importance of that dimension to California players (see Table 1.2).
<b>Sampling fraction</b>	The ratio of the study sample to the estimated active population. Here approximately one percent, computed against the estimated active California online-casino audience of roughly 420,000 rather than against the general state population.
<b>Funded-account screening</b>	The verification procedure requiring that a respondent must have funded a real-money account at a named casino within the preceding twelve months before that respondent's answers may contribute to that casino's California Player Survey sub-score.
<b>Verified respondent</b>	A participant who passed email and SMS verification, government-identification age verification confirming age twenty-one or older, device-fingerprint de-duplication, and the embedded attention-check items.
<b>Limited player data</b>	A flag applied to any operator that did not reach the minimum of one hundred verified funded respondents and is therefore assessed on tester evidence rather than a survey-based California Player Survey sub-score.
<b>Grey market</b>	A market in which offshore operators serve players in a jurisdiction that neither licenses nor prohibits the activity outright, leaving players without domestic regulatory recourse. California's online-casino market is treated here as a grey market.
<b>Conventional star rating</b>	A generic or national operator rating, typically on a five-star scale, produced without California-specific calibration or funded California player verification, used in this study as a comparison benchmark for the CAC Score.

Table 1.2 summarizes the eight pillars and their respective weights, which together govern the computation of every CAC Score reported in this thesis. The weights are introduced here for reference and are justified in detail in the methodology chapter.

**Table 1.2 — The eight CAC Score pillars and their weights (summing to 100)**

Pillar	Weight (%)
California Player Survey	20
Trust & Licensing	18
Payout Speed & Banking	15
Bonuses & Value	14
Game Selection	13
Security & Fairness	8
Customer Support	7
Mobile & Responsible Gambling	5

*Figure 1.1 — Relative pillar weights in the CAC Score composite, illustrating the dominance of the California Player Survey and Trust & Licensing pillars.*

## 1.8 Thesis Structure

The remainder of this thesis is organized into six further chapters, each building on the foundations laid in the present introduction.

**Chapter 2 (Literature Review)** situates the study within the existing scholarship on online gambling behaviour, player protection in unregulated and grey markets, and the methodology of consumer-facing rating systems. It reviews prevalence research on online gambling, the literature on information asymmetry between operators and players, and prior attempts to construct composite quality indices, identifying the gap that the California-specific framework is designed to fill.

**Chapter 3 (Methodology)** sets out the full research design. It details the population and sampling frame, the stratified random and snowball recruitment mix, the regional and demographic quotas, the four-stage verification and screening protocol, the sixty-four-item eight-domain instrument, the funded-account rule, and the construction and weighting of the eight pillars. It also specifies the statistical procedures employed, including descriptive statistics, independent-samples t-tests, one-way analysis of variance across regions, confidence-interval and margin-of-error computation, Cohen's d effect sizes, and Cronbach's alpha reliability analysis.

**Chapter 4 (Results)** reports the empirical findings. It presents the sample composition against the target quotas, the reliability statistics for the instrument, the pillar sub-scores and composite CAC Scores for each evaluated operator, the distribution of operators across performance bands, and the inferential tests addressing the study's hypotheses.

**Chapter 5 (Discussion)** interprets these findings in light of the research questions. It examines the divergence between CAC Scores and conventional star ratings, considers the player-protection implications of the observed market distribution, and reflects on what the funded-account rule reveals about the reliability of unverified rating systems.

**Chapter 6 (Limitations and Future Research)** appraises the boundaries of the study, including the cross-sectional fieldwork window, the reliance on online self-administration, and the constraints imposed by the verification thresholds, and it outlines directions for longitudinal extension and cross-jurisdictional replication.

**Chapter 7 (Conclusion)** synthesizes the contributions of the study, restates the answers to the research questions, and articulates the practical recommendations for California players, player-protection advocates, and future researchers. The References chapter and appendices follow, compiling all cited sources in APA seventh-edition format and documenting the instrument and supporting materials.

## Chapter 2 — Literature Review

---

### 2.1 Introduction to the review

This chapter establishes the theoretical and empirical foundations on which the California Casino Score (CAC Score) framework rests. The aim of the review is twofold. First, it situates the present study within the wider body of scholarship on online gambling behaviour, survey methodology, psychometric measurement, and gambling regulation, demonstrating that the study's design decisions are grounded in established theory rather than convenience. Second, it interrogates the existing landscape of consumer-facing casino rating and review systems, identifying the methodological weaknesses that the CAC Score is explicitly designed to remedy. The review therefore functions both as a justification for the eight-pillar weighted instrument and as a map of the gap the study claims to fill.

The literature is approached thematically rather than chronologically, on the basis that the research problem is inherently multidisciplinary. Evaluating offshore online casinos for a regional population draws simultaneously on behavioural economics, consumer psychology, survey statistics, measurement theory, and public policy. A purely chronological treatment would obscure the way these strands intersect at the point of instrument design. Accordingly, the chapter proceeds from the broad behavioural context (Section 2.2), through the critique of existing rating methodologies (Section 2.3), into the methodological machinery the study employs (Sections 2.4 and 2.5), then to the substantive constructs underlying the eight pillars (Section 2.6), the regulatory context specific to California (Section 2.7), and finally to a synthesis of the identified gaps (Section 2.8) and the conceptual framework that links theory to the weighted scoring model (Section 2.9).

Throughout, sources are cited in APA 7th edition author–date format, consistent with the conventions adopted across the thesis. Where the literature is contested, competing positions are presented rather than a single view being asserted. The review draws on peer-reviewed gambling-studies journals, methodological texts, regulatory publications, and the grey literature of consumer-protection bodies, reflecting the applied nature of the research problem.

#### 2.1.1 Scope and boundaries

The review is bounded in three respects. Geographically, it privileges literature relevant to North American and, where available, Californian gambling markets, while drawing on international evidence (notably from the United Kingdom, Australia, and continental Europe) where domestic studies are sparse. Substantively, it concentrates on real-money online casino play rather than land-based gambling, sports

betting in isolation, or social-casino products, although these adjacent literatures are referenced where they illuminate the focal construct. Methodologically, it foregrounds survey-based and mixed-methods research, since the CAC Score combines a large stratified survey with structured expert testing. Lottery and state-monopoly contexts are treated only insofar as they inform the regulatory discussion in Section 2.7.

## **2.2 Online gambling research and consumer behaviour**

Online gambling has been one of the fastest-growing segments of the global gambling industry over the past two decades, and a substantial research literature has accumulated around questions of adoption, motivation, trust, and harm (Gainsbury, 2015; Hing et al., 2014). The migration of gambling from land-based venues to internet and, increasingly, mobile platforms has altered both the structural characteristics of products and the behavioural profile of the players who engage with them. This section reviews four interlocking themes: adoption and channel choice, trust as a precondition for participation, player motivations, and gambling-related harm.

### ***2.2.1 Adoption and the shift to mobile***

Research on the adoption of online gambling has frequently been framed through technology-acceptance and diffusion models, in which perceived usefulness, perceived ease of use, and perceived risk jointly shape intention to use a platform (Davis, 1989; Rogers, 2003). Studies adapting these models to gambling contexts report that convenience and continuous availability are primary adoption drivers, while concerns about payment security and operator legitimacy act as the principal brakes (Gainsbury et al., 2013). The accelerating dominance of mobile play is well documented: smartphones have become the primary gambling device for a majority of online players in several jurisdictions, a pattern strongly reflected in the present study, where mobile is the primary device for 64% of respondents, against 28% for desktop and 8% for tablet. This distribution is consistent with international findings that mobile-first behaviour is now the modal pattern among younger cohorts in particular (James, O'Malley, & Tunney, 2017). The behavioural significance of mobile play—shorter, more frequent sessions and heightened exposure to push-based marketing—has direct implications for how mobile experience and responsible-gambling features are weighted in any evaluative instrument.

### ***2.2.2 Trust as a precondition for participation***

In an online environment where the player cannot physically inspect the operator, trust functions as a precondition for participation rather than as a downstream outcome (Gefen, Karahanna, & Straub, 2003).

The e-commerce trust literature distinguishes between institution-based trust (reliance on regulatory and licensing structures), calculative trust (assessment of the operator's incentives), and experiential trust accumulated through repeated transactions (McKnight, Choudhury, & Kacmar, 2002). Applied to gambling, institution-based trust is heavily mediated by licensing signals, while experiential trust is dominated by the single most consequential transaction in the player journey: the withdrawal. Empirical work indicates that perceived payout reliability and the visible presence of a recognised licence are among the strongest predictors of willingness to deposit and to recommend an operator (Wohl et al., 2017). This evidence directly motivates the prominence of the Trust & Licensing and Payout Speed & Banking pillars in the CAC Score, which together carry a third of the total weight.

### ***2.2.3 Player motivations***

Motivational research consistently identifies a heterogeneous set of drivers for gambling, including monetary aspiration, entertainment and escape, social interaction, and the intrinsic appeal of risk and challenge (Lee et al., 2007; Binde, 2013). Segmentation studies show that motivational profiles vary systematically by age, gender, and product preference, with slots players often emphasising entertainment and immersion, and table or live-dealer players more frequently citing skill perception and social presence (Hing et al., 2014). The implication for evaluative frameworks is that a single undifferentiated quality score risks averaging away the very dimensions that distinct player segments value most. The CAC Score's decomposition into game selection, bonuses and value, support, and player-reported experience reflects this motivational heterogeneity rather than collapsing it.

### ***2.2.4 Gambling-related harm***

A parallel and ethically central literature addresses gambling-related harm. Online gambling has been associated in some studies with elevated risk of disordered play relative to land-based gambling, attributable to accessibility, immersion, and the speed of play cycles (Gainsbury et al., 2015). Public-health framings increasingly treat harm as a continuum experienced across a population rather than as a binary condition affecting a clinical minority (Browne et al., 2016). This reframing has methodological consequences: it argues for incorporating responsible-gambling provision into any quality assessment and for surfacing harm-minimisation resources to research participants themselves. The present study operationalises both commitments, weighting Mobile & Responsible Gambling within the scoring model and surfacing BeGambleAware and the National Council on Problem Gambling helpline (1-800-522-4700) to all participants, consistent with the ethical norms articulated in the responsible-gambling literature (Ladouceur et al., 2017).

## 2.3 Existing casino rating and review methodologies and their limitations

The consumer-facing casino review sector is large, commercially driven, and methodologically under-scrutinised. While academic gambling research is extensive, the rating systems that actually shape player decisions—star ratings, "top ten" lists, and editorial scores published by affiliate websites—have received comparatively little independent evaluation (Marionneau & Nikkinen, 2020). This section identifies four recurrent limitations of the prevailing methodologies and motivates the corrective design choices embedded in the CAC Score.

### 2.3.1 Generic star ratings and construct opacity

The dominant rating idiom is the undifferentiated star score, typically expressed on a five-point scale. Such ratings are attractive for their cognitive simplicity but suffer from acute construct opacity: the dimensions being aggregated, their relative weights, and the evidence base behind each are rarely disclosed (Marionneau & Nikkinen, 2020). A measurement-theory critique holds that a score whose underlying construct and weighting scheme are unspecified cannot be evaluated for validity, because there is no stated referent against which to judge it (Messick, 1995). The empirical inadequacy of the generic star score is illustrated within the present study's own dataset. Table 2.1 contrasts the published five-point star rating of the evaluated casinos with their multidimensional CAC Score. The two measures are broadly monotonic, but the star scale *compresses* information that the CAC Score preserves: a 23-point CAC range (from Ignition at 98 down to VoltageBet at 75) is collapsed onto only five discrete star steps. The 3.5-star band alone spans four CAC points (Wild 86 down to Lucky Creek 82), so a star-only reader treats as equivalent operators that the decomposed measure separates by a margin equal to a full half-star step. A single-number star rating thus conceals far more gradation than it reveals.

**Table 2.1 — Information loss: the five-point star rating versus the multidimensional CAC Score (study dataset, sorted by CAC Score)**

Casino	Published star rating (0–5)	CAC Score (0–100)	Verified CA players
Ignition	5.0	98	107
BetOnline	5.0	97	249
All Star Slots	4.5	96	194
Super Slots	4.5	95	212
Slots.lv	4.0	91	126
Slots of Vegas	4.0	90	222
Bovada	4.0	89	123
Wild	3.5	86	203
Lucky Creek	3.5	82	230
Shazam	3.0	77	216
VoltageBet	3.0	75	106

The pattern in Table 2.1 demonstrates that while the star rating tracks the direction of quality, it discards resolution: within every star band the CAC Score varies by several points, and adjacent bands abut so closely that a single point of CAC perception can move a casino across a half-star boundary. A coarse headline score is therefore an adequate device for triage but a poor proxy for ranking the dimensions that the behavioural literature identifies as decision-relevant (Wohl et al., 2017).

### **2.3.2 Affiliate bias and commercial conflict of interest**

A second and widely noted limitation is structural conflict of interest. The majority of high-traffic casino review sites operate on affiliate-marketing models, earning commission on player referrals and deposits (Newall et al., 2019). This commercial dependency creates an incentive to inflate ratings of high-commission operators and to suppress critical findings, a dynamic documented across affiliate-driven sectors more broadly (Mathur, Narayanan, & Chetty, 2018). The gambling-specific consequence is that published rankings may reflect commercial relationships rather than product quality or player welfare (Newall et al., 2019). Independence from operators is therefore a prerequisite for methodological credibility, a principle the present study adopts explicitly by maintaining no operator affiliation and disclosing this position to participants and readers.

### **2.3.3 Absence of regional specificity**

A third limitation is the near-universal absence of regional specificity. Most rating systems publish a single global score that ignores the substantial jurisdictional variation in payment rails, currency support, regulatory exposure, customer-support availability, and product legality (Gainsbury, 2012). A casino that

performs well for European players may serve North American players poorly on withdrawal speed, banking options, or support hours. The behavioural literature on payout reliability (Section 2.2.2) implies that such regionally contingent factors are precisely those that most strongly shape trust and retention, yet they are flattened by global scoring. The CAC Score's defining innovation is to anchor evaluation to a single, well-defined regional population—California residents aged 21 and over who played real-money online casino games in the preceding twelve months—thereby restoring the regional specificity that generic systems discard.

### 2.3.4 Lack of player data

The fourth and arguably most consequential limitation is the scarcity of genuine player data underpinning published ratings. Editorial scores are predominantly expert-derived or marketing-derived; few are grounded in verified evidence from real depositors at the named operator (Marionneau & Nikkinen, 2020). Where player sentiment is incorporated, it is typically drawn from unverified open review platforms vulnerable to manipulation, fake reviews, and sampling bias (Luca & Zervas, 2016). The present study addresses this directly through a validity rule: a casino receives a survey-based California Player Survey sub-score only where at least 100 verified respondents have funded a real-money account at that operator within the previous twelve months; casinos below this threshold are tester-assessed and explicitly flagged as having "limited player data". This thresholding converts player data from a marketing veneer into a stated evidentiary requirement.

**Table 2.2 — Comparison of rating methodologies against four limitation criteria**

Methodology	Construct transparency	Independence from operators	Regional specificity	Verified player data
Generic star rating	Low	Often compromised	None	Rare / unverified
Affiliate "top list"	Low	Low (commission-driven)	None	Marketing-derived
Open user-review platform	Moderate	Variable	Limited	Unverified, manipulable
Expert editorial score	Moderate	Variable	Limited	Expert, not player
CAC Score (this study)	High (8 disclosed weighted pillars)	Stated independence	High (CA 21+ population)	Verified, threshold-gated

## 2.4 Survey methodology and sampling theory

Because the CAC Score's most heavily weighted pillar is the California Player Survey, the credibility of the entire framework depends on the soundness of its sampling and survey design. This section reviews the relevant methodological theory and shows how each principle is operationalised in the study.

### ***2.4.1 Probability versus non-probability sampling***

A foundational distinction in survey research separates probability sampling, in which every population element has a known, non-zero selection probability, from non-probability sampling, in which selection rests on accessibility or judgement (Lohr, 2019; Groves et al., 2009). Probability methods permit defensible inference to the target population and the calculation of sampling error; non-probability methods sacrifice this in exchange for feasibility, particularly when the population is hidden or hard to enumerate (Baker et al., 2013). Online gambling populations are difficult to sample probabilistically because no exhaustive sampling frame of real-money players exists. The present study adopts a pragmatic hybrid in response: 78% of respondents are drawn from a stratified random panel and 22% through snowball referral. This mix is defended in the methodological literature as a reasonable compromise for hidden populations, provided the non-probability component is bounded and disclosed (Heckathorn, 1997).

### ***2.4.2 Stratified sampling***

Stratified sampling partitions the population into mutually exclusive strata and samples within each, typically improving precision relative to simple random sampling when strata are internally homogeneous and externally heterogeneous (Lohr, 2019). Proportional allocation, in which each stratum's sample share matches its population share, preserves representativeness on the stratifying variable (Cochran, 1977). The present study stratifies the Californian player population by region using proportional quotas: Southern California 58%, San Francisco Bay Area 20%, Central Valley 11%, Sacramento Metro 6%, Central Coast 3%, and North State 2%. This regional stratification is what enables the one-way analysis of variance across regions reported in later chapters, and it operationalises the regional-specificity principle established in Section 2.3.3. Secondary quotas on age band (21–29: 22%; 30–44: 38%; 45–59: 27%; 60+: 13%) and gender (men 56%; women 43%; nonbinary/undisclosed 1%) further align the achieved sample with the assumed population profile.

### ***2.4.3 Snowball sampling for hidden populations***

Snowball and respondent-driven sampling are established techniques for reaching populations that lack a sampling frame, relying on existing respondents to refer further participants (Goodman, 1961; Heckathorn, 1997). Their well-documented weakness is dependence on social-network structure, which can bias the sample toward connected subgroups (Baker et al., 2013). The mitigation adopted here is to confine snowball recruitment to a minority share (22%) of the total and to subject all referred respondents

to the same verification and screening regime as panel respondents, limiting the influence of referral chains on the aggregate estimates.

#### **2.4.4 Sample size, sampling fraction, and margin of error**

The precision of a survey estimate is governed by the relationship between sample size, population variability, and the desired confidence level (Cochran, 1977). For a proportion, the margin of error at a given confidence level is approximately a function of the estimated proportion and the square root of the sample size, with a finite-population correction applied when the sampling fraction is non-trivial (Lohr, 2019). The present study recruits a final, post-screening sample of  $N = 4,217$  verified respondents from an estimated active California online-casino audience of approximately 420,000, yielding a sampling fraction of roughly 1% of that active audience—not 1% of the general California population. At this sample size the study reports a 95% confidence level with a margin of error of approximately  $\pm 1.5\%$ . Figure 2.1 illustrates the conventional inverse-square-root relationship between sample size and margin of error that underlies this design, clarifying why increments beyond several thousand respondents yield diminishing precision gains.

*Figure 2.1 — Schematic relationship between sample size and margin of error at 95% confidence, showing diminishing precision returns and the position of the study's achieved sample ( $N = 4,217$ ,  $MoE \approx \pm 1.5\%$ ).*

#### **2.4.5 Verification and screening as data-quality controls**

Online survey research is vulnerable to fraudulent, duplicate, and inattentive responses, which inflate sample size while degrading data quality (Chandler & Paolacci, 2017). The methodological literature recommends layered controls: identity verification, de-duplication, attention checks, and screening on eligibility criteria. The present study implements email and SMS verification, government-identification age verification to enforce the 21-and-over criterion, device-fingerprint de-duplication, and three embedded attention-check items, with casino-specific items further gated on funded-account screening. These controls are what permit the description of the sample as "verified" rather than merely "responding", and they directly support the validity threshold described in Section 2.3.4.

### **2.5 Measurement: Likert scales, instrument design, validity and reliability**

Once a sample is secured, the quality of inference depends on the measurement properties of the instrument. This section reviews the psychometric theory governing the study's questionnaire.

#### **2.5.1 Likert scaling and attitudinal measurement**

The Likert scale remains the dominant format for measuring attitudes through summated agreement ratings (Likert, 1932). Debate persists over whether Likert data should be treated as ordinal or interval, with the pragmatic position holding that summated multi-item scales approximate interval measurement adequately for parametric analysis (Carifio & Perla, 2008). The choice of scale length is also consequential: longer scales (seven points) can improve discrimination and reliability relative to shorter ones without unduly burdening respondents (Preston & Colman, 2000). The present study employs seven-point Likert agreement scales for attitudinal items, alongside behavioural and factual items used for verification, distributing 64 items across the eight domains. This design reflects the recommendation that attitudinal constructs be measured with multiple items per construct to support reliability estimation (DeVellis, 2017).

### ***2.5.2 Instrument design and content validity***

Sound instrument design begins with a clearly specified construct map and proceeds through item generation, expert review, and pilot testing to establish content validity—the degree to which items adequately sample the construct domain (DeVellis, 2017; Messick, 1995). The eight-domain structure of the questionnaire mirrors the eight pillars of the scoring model, so that each construct in the conceptual framework is represented by a dedicated block of items. This one-to-one mapping between theoretical construct and measured domain is a deliberate content-validity strategy, ensuring that the instrument samples the full breadth of the construct space the framework claims to assess.

### ***2.5.3 Reliability and Cronbach's alpha***

Reliability concerns the consistency of measurement, with internal-consistency reliability—the extent to which items within a scale measure a common construct—most commonly estimated by Cronbach's alpha (Cronbach, 1951). Conventional benchmarks treat alpha values of 0.70 and above as acceptable, 0.80 and above as good, and 0.90 and above as excellent, while cautioning that very high values can signal item redundancy (Nunnally & Bernstein, 1994; Tavakol & Dennick, 2011). The present study reports an overall internal-consistency reliability of Cronbach's alpha = 0.89, with domain-level alphas ranging from 0.78 to 0.92. These values fall within the conventionally acceptable-to-excellent range, supporting the claim that the multi-item domains measure coherent underlying constructs. Table 2.3 maps the standard interpretive bands to the study's reported figures.

**Table 2.3 — Conventional Cronbach's alpha interpretation bands and the study's reported reliability**

Alpha range	Conventional interpretation	Study reference point
≥ 0.90	Excellent (watch for redundancy)	Upper domain alpha = 0.92
0.80–0.89	Good	Overall alpha = 0.89
0.70–0.79	Acceptable	Lower domain alpha = 0.78
< 0.70	Questionable to unacceptable	None reported

### **2.5.4 Validity, effect sizes, and inferential statistics**

Beyond reliability, the validity of inferences drawn from an instrument depends on appropriate analytical procedures (Messick, 1995). The study's analytical plan draws on a standard inferential toolkit: descriptive statistics (means, standard deviations, frequencies), independent-samples t-tests, one-way ANOVA across the regional strata, 95% confidence intervals, and margin-of-error computation. Crucially, it also reports Cohen's d effect sizes, addressing the widely voiced criticism that statistical significance alone, particularly with large samples, can mislead by flagging trivial differences as "significant" (Cohen, 1988; Cumming, 2014). Reporting effect sizes alongside significance tests is consistent with contemporary best-practice recommendations for transparent quantitative reporting (Cumming, 2014).

## **2.6 Trust, licensing, payout and player-protection literature**

This section reviews the substantive constructs underlying the non-survey pillars of the CAC Score, establishing that each weighted dimension corresponds to a documented determinant of player welfare or decision-making.

### **2.6.1 Licensing and institutional trust**

Licensing operates as the principal institutional trust signal in online gambling, certifying that an operator is subject to external oversight regarding fairness, segregation of player funds, and dispute resolution (Gainsbury, 2012). However, the offshore casinos most accessible to United States players are frequently licensed in jurisdictions with lighter regulatory regimes, so the strength of the licensing signal varies considerably (Marionneau & Nikkinen, 2020). The literature's emphasis on institution-based trust (Section 2.2.2) justifies assigning Trust & Licensing the second-largest weight in the model at 18%.

### **2.6.2 Payout speed and banking**

Withdrawal experience is repeatedly identified as the decisive moment in the player–operator relationship, with delayed, denied, or friction-laden payouts being among the most common and damaging player grievances (Wohl et al., 2017). The rise of cryptocurrency banking has reshaped this dimension, offering faster settlement but introducing volatility and usability barriers for some segments

(Gainsbury & Blaszczynski, 2017). The Payout Speed & Banking pillar, weighted at 15%, operationalises this literature.

### ***2.6.3 Bonuses, value, and the role of terms***

Promotional bonuses are central to acquisition and retention but are frequently encumbered by wagering requirements and conditions that erode their realisable value (Newall et al., 2019). The consumer-protection literature stresses transparency of terms as the relevant quality criterion rather than headline bonus size (Newall et al., 2019). The Bonuses & Value pillar (14%) reflects this value-not-volume orientation.

### ***2.6.4 Game selection, security, fairness, and support***

Game breadth and provider quality drive the entertainment motivations documented in Section 2.2.3 and inform the Game Selection pillar (13%). Security and fairness—encompassing data protection, encryption, and the use of certified random number generators and audited return-to-player percentages—underpin both trust and harm-related concerns (Gainsbury, 2012) and are captured in the Security & Fairness pillar (8%). Customer support quality, an experiential determinant of trust repair when problems arise, forms a dedicated pillar (7%).

### ***2.6.5 Responsible gambling and player protection***

The player-protection literature, building on the public-health framing of Section 2.2.4, specifies a core set of harm-minimisation features—deposit limits, self-exclusion, reality checks, and visible signposting to help resources—as markers of operator responsibility (Ladouceur et al., 2017; Browne et al., 2016). These features, together with mobile usability, constitute the Mobile & Responsible Gambling pillar (5%). Although the smallest weight in the model, its inclusion signals that welfare considerations are embedded in the evaluative construct rather than excluded from it.

## **2.7 California regulatory and market context**

The framework is regionally bounded, so its interpretation requires an account of the Californian regulatory and market environment. This section reviews the relevant policy literature.

### ***2.7.1 Tribal gaming and the legal landscape***

California's land-based gambling market is structured around tribal gaming, conducted under the federal Indian Gaming Regulatory Act framework and tribal–state compacts, alongside licensed card rooms and the state lottery (Rose, 2020). Tribal interests have historically been influential in shaping the state's

gambling policy, including the contours of any prospective expansion (Rose, 2020). Crucially, the state has not legalised and licensed online casino play, leaving a regulatory vacuum at the internet-casino level even as adjacent verticals are debated.

### ***2.7.2 The offshore market and its risks***

In the absence of a domestic licensing regime for online casinos, Californian players who wish to gamble online predominantly access offshore operators licensed in foreign jurisdictions (Gainsbury, 2012; Marionneau & Nikkinen, 2020). This offshore reality has direct consequences for the constructs the framework measures: regulatory recourse is limited, payout enforcement rests largely on operator reputation, and player-protection standards are not domestically mandated. The CAC Score's emphasis on verified player experience and disclosed licensing is a direct response to this comparatively unprotected market structure.

### ***2.7.3 Ballot measures and the 21-plus norm***

California's gambling policy has repeatedly been contested through the ballot-initiative process, where competing tribal, commercial, and operator interests have advanced rival expansion proposals (Rose, 2020). For the purposes of this study, the salient regulatory norm is the minimum gambling age. Consistent with the prevailing 21-and-over standard applied to much commercial and tribal gambling in the state, the study restricts its population to residents aged 21 and over and enforces this through government-identification age verification, a decision both methodologically and ethically grounded in the Californian context.

## **2.8 Identified gaps in the literature**

Synthesising the preceding sections reveals a coherent set of gaps that the present study is designed to address. Four are foremost.

First, there is a measurement-transparency gap. The behavioural literature richly documents the determinants of player trust and choice (Section 2.2), yet the rating systems that actually guide players obscure their constructs and weights (Section 2.3.1). No widely used framework publishes a fully disclosed, theory-aligned weighting scheme of the kind the CAC Score sets out.

Second, there is an independence gap. The dominance of affiliate-funded rating systems (Section 2.3.2) means that few publicly available evaluations are produced under a stated independence-from-operators commitment, leaving consumers without a credibly disinterested benchmark.

Third, there is a regional-specificity gap. The literature establishes that payout, banking, support, and legality are jurisdictionally contingent (Sections 2.3.3 and 2.7), yet global single-score systems systematically ignore this, and no prior framework anchors itself to the specific construct of the Californian 21-plus online-casino player.

Fourth, there is a verified-player-data gap. Existing systems rely on expert or marketing inputs, or on manipulable open reviews (Sections 2.3.4 and 2.5.5), with no enforced evidentiary threshold tying a player-experience score to a minimum count of verified depositors. The study's 100-verified-funder rule has no documented precedent in the consumer casino-rating literature.

These four gaps—transparency, independence, regional specificity, and verified data—are not independent; they compound one another, since an opaque, conflicted, global, marketing-sourced score fails consumers on every axis simultaneously. The CAC Score is constructed to close all four within a single, internally consistent framework.

## **2.9 Conceptual framework**

The conceptual framework draws the preceding literature together into the operational logic of the CAC Score. Its central proposition is that the overall quality of an offshore online casino, as experienced by a Californian player, is a weighted composite of eight theoretically grounded dimensions, each measured through a combination of verified player report and structured expert testing, and each weighted in proportion to its documented importance to player welfare and decision-making.

### ***2.9.1 Linking the eight pillars to the literature***

Each pillar is anchored to a specific body of evidence reviewed above. The California Player Survey pillar, carrying the largest weight at 20%, instantiates the verified-player-data principle (Sections 2.3.4 and 2.4.5) and the motivational-heterogeneity argument (Section 2.2.3). Trust & Licensing (18%) operationalises institution-based trust (Sections 2.2.2 and 2.6.1). Payout Speed & Banking (15%) reflects the decisive role of withdrawal experience (Section 2.6.2). Bonuses & Value (14%) encodes the transparency-of-terms criterion (Section 2.6.3). Game Selection (13%) maps to entertainment motivations and product breadth (Sections 2.2.3 and 2.6.4). Security & Fairness (8%) draws on the fairness and data-protection literature (Section 2.6.4). Customer Support (7%) captures experiential trust repair (Section 2.6.4). Mobile & Responsible Gambling (5%) embeds the public-health and harm-minimisation literature (Sections 2.2.4 and 2.6.5). The weights, which sum to 100, thus encode a literature-justified hierarchy of importance rather than an arbitrary allocation.

*Figure 2.2 — Conceptual framework of the CAC Score: eight literature-anchored pillars (weights summing to 100) feeding a single 0–100 composite, with verified-player gating on the survey pillar and effect-size-reported regional analysis.*

### **2.9.2 The scoring logic and its banded interpretation**

The framework aggregates the eight pillar sub-scores, each expressed on a 0–100 scale, into a single composite computed as the weight-multiplied sum divided by 100, reported on a 0–100 scale and interpreted through defined bands (85 and above Excellent, 75–84 Good, 65–74 Fair, below 65 Weak). This compensatory weighted-additive structure is the standard form for multi-criteria evaluative composites and is defensible where the dimensions are conceptually distinct and individually interpretable (Messick, 1995), conditions the reliability evidence in Section 2.5.3 supports. The validity gate—awarding a survey-based player sub-score only where at least 100 verified funders exist—is the framework's mechanism for ensuring that the most heavily weighted pillar rests on adequate evidence, with under-threshold casinos transparently flagged as having limited player data.

### **2.9.3 Summary and transition to method**

The conceptual framework therefore translates the gaps identified in Section 2.8 into a concrete, measurable, and weighted model, with each component traceable to the literature reviewed in this chapter. Transparency is achieved through fully disclosed pillars and weights; independence through a stated non-affiliation commitment; regional specificity through a population defined as Californian residents aged 21 and over; and verified-player grounding through layered verification and the funded-account threshold. The chapter that follows operationalises this framework into the detailed research design, sampling protocol, and analytical procedures of the study.

## Chapter 3 — Methodology

---

This chapter sets out the methodological framework underpinning the California Casino Score (CAC Score) study, a stratified survey-and-testing programme conducted by the Palo Alto Casino Analysis Collective (operating as "CA Casinos", [cacpaloalto.org](http://cacpaloalto.org)) between January and April 2026. The chapter describes the research design, the target population and sampling frame, the sampling strategy and its statistical justification, the recruitment and verification procedures, the funded-account screening rule that governs which casinos may receive a survey-based sub-score, the design of the measurement instrument, the reliability and validity evidence assembled for that instrument, the weighting model that converts pillar sub-scores into a composite CAC Score, the planned statistical analyses, and the ethical safeguards applied throughout. The intention is to provide sufficient procedural detail for the study to be evaluated, replicated, and extended by independent researchers.

### 3.1 Research design

The study adopts a convergent mixed-methods design in which two complementary streams of evidence are collected in parallel and integrated at the point of scoring. The first stream is a large-scale, self-administered online survey of California residents who play real-money online casino games. The second stream is a programme of structured operator testing, in which trained analysts conduct standardised assessments of each casino across attributes that players cannot reliably self-report, such as withdrawal processing times, banking-method availability, licence verification, and the technical implementation of security and fairness controls.

The rationale for combining the two streams is methodological triangulation (Creswell & Plano Clark, 2018). Player-reported attitudes capture lived experience, perceived value, and satisfaction, but they are vulnerable to recall bias, exposure bias, and uneven sample coverage across operators. Structured testing, by contrast, produces reproducible, operator-independent observations but cannot speak to subjective player experience. By weighting and aggregating both streams within a single composite index, the design preserves the explanatory richness of self-report while anchoring the final score to verifiable operational facts. The survey stream is quantitative and cross-sectional; the testing stream is a standardised observational protocol. Neither stream is treated as subordinate; rather, each pillar of the final index is sourced from whichever stream offers the more valid measurement for the attribute in question, as detailed in Section 3.6.

The unit of analysis is dual. For population-level inference (for example, regional differences in player satisfaction), the unit is the individual verified respondent. For the construction and reporting of the CAC Score, the unit is the individual casino. This dual structure is reflected throughout the analysis plan in Section 3.9.

A cross-sectional rather than longitudinal survey design was adopted for the player stream. The research questions concern the current state of the California offshore-casino market and the relative standing of operators at a defined point in time, rather than the trajectory of individual players' attitudes. A cross-sectional design answers those questions efficiently and avoids the panel attrition, conditioning, and re-identification risks that a longitudinal design would introduce in a population for whom privacy is an acute concern. Future waves of the programme are expected to repeat the cross-section, enabling between-wave comparison at the market level without retaining individual-level tracking data, which preserves the data-minimisation commitments described in Section 3.10.

The two streams differ in their epistemological footing, and the design treats that difference deliberately. The survey stream yields probabilistic, population-referenced estimates whose uncertainty is quantified through confidence intervals and margins of error. The testing stream yields deterministic, protocol-referenced observations whose credibility rests on the standardisation and repeatability of the testing procedure rather than on sampling theory. Integrating two evidence types with different uncertainty structures requires that each pillar declare its provenance, so that a reader can interpret a pillar sub-score in light of how it was produced. This provenance is made explicit through the funded-account flag (Section 3.5) and through the pillar-by-pillar mapping in Section 3.6.

*Figure 3.1 — Convergent mixed-methods design: parallel survey and structured-testing streams integrated at the pillar-scoring stage to produce a single composite CAC Score per casino.*

## **3.2 Population and sampling frame**

The target population is defined as California residents aged 21 years or older who played real-money online casino games at least once in the twelve months preceding fieldwork. The minimum-age criterion of 21 reflects the most conservative threshold applicable to the relevant gambling activities and aligns the study with the operator-facing screening used in the verification procedure (Section 3.4). Residents below 21, residents who had not played real-money online casino games in the reference period, and individuals residing outside California were outside the population of interest and were screened out prior to enrolment.

The estimated size of the active California online-casino audience is approximately 420,000 players. This figure constitutes the sampling frame against which the achieved sample is benchmarked and from which the sampling fraction is computed. It is important to emphasise that the frame is the active online-casino-playing audience, not the general adult population of California; the sampling fraction reported in Section 3.3 is therefore expressed relative to that 420,000-player audience and not relative to the tens of millions of California residents overall. Anchoring the frame to active players rather than the general population improves the relevance of the recruited sample to the constructs under study, because every eligible respondent has direct, recent experience of the behaviours being measured.

Fieldwork was conducted over a four-month window (January–April 2026). A multi-month window was chosen to dampen short-term seasonality in play behaviour (for example, promotional cycles tied to major sporting events) and to allow sufficient time for the layered verification and screening procedures to run without compromising data quality in pursuit of speed.

### **3.3 Sampling strategy**

The study employed a dual-frame sampling strategy combining stratified random sampling with snowball referral. Of the achieved sample, 78% was recruited through a stratified random panel and 22% through snowball referral. The stratified random component is the methodological backbone of the design and supports the inferential claims made in Chapter 4; the snowball component was incorporated to reach lower-incidence segments of the player population (for example, players of niche operators and players in sparsely populated regions) that probability panels tend to under-cover (Heckathorn, 1997).

#### ***3.3.1 Stratification and regional quotas***

Stratification was applied on region of residence, with proportional quotas designed to mirror the regional distribution of the active California player audience. Six regional strata were used. Within each stratum, panel members meeting the eligibility criteria were selected at random until the proportional quota was filled. The quotas are reported in Table 3.1.

**Table 3.1 — Regional strata and proportional sampling quotas**

Regional stratum	Quota (% of sample)	Approx. respondents (of N = 4,217)
Southern California	58%	2,446
SF Bay Area	20%	843
Central Valley	11%	464
Sacramento Metro	6%	253
Central Coast	3%	127
North State	2%	84
<b>Total</b>	<b>100%</b>	<b>4,217</b>

Secondary monitoring quotas were maintained on age band, gender, and primary device. These secondary distributions were tracked during fieldwork and used to identify and correct under-representation through targeted further recruitment, but the primary stratification controlling random selection was regional. The achieved secondary distributions are reported in Table 3.1b. Monitoring these characteristics serves two purposes: it allows the sample composition to be benchmarked against the known profile of the active player audience, and it provides the grouping variables for the sub-group comparisons set out in the analysis plan (Section 3.9).

**Table 3.1b — Secondary monitoring quotas (age band, gender, primary device)**

Characteristic	Category	Share of sample
Age band	21–29	22%
	30–44	38%
	45–59	27%
	60+	13%
Gender	Men	56%
	Women	43%
	Nonbinary / undisclosed	1%
Primary device	Mobile	64%
	Desktop	28%
	Tablet	8%

The dual-frame design carries a known methodological tension that the study acknowledges explicitly. The stratified random component supports unbiased population inference, whereas the snowball component, being a non-probability referral method, can over-represent socially connected players and the operators they favour. To contain this, the snowball component was capped at 22% of the achieved sample, the two components were tagged at the record level so that frame membership could be examined during analysis, and the regional quotas were enforced across the combined sample so that referral chains

could not distort the geographic distribution. Where a sub-group estimate is materially sensitive to frame membership, that sensitivity is reported in Chapter 4 rather than suppressed.

### 3.3.2 Sample size and sampling fraction

The final achieved sample, after all verification and screening steps described in Sections 3.4 and 3.5, comprised  $N = 4,217$  verified respondents. Relative to the estimated active audience of approximately 420,000 players, this corresponds to a sampling fraction of approximately 1%, computed as  $4,217 / 420,000 = 0.01004$ , that is, 1.00%.

### 3.3.3 Confidence level and margin of error

The study reports estimates at the 95% confidence level with a margin of error of  $\pm 1.5\%$  for proportions estimated on the full sample. The margin of error for a proportion is computed using the standard large-sample formula:

$$\text{MoE} = z \times \sqrt{(p(1 - p) / n)}$$

where  $z$  is the critical value of the standard normal distribution for the chosen confidence level ( $z = 1.96$  for 95% confidence),  $p$  is the sample proportion, and  $n$  is the sample size. Following standard practice, the most conservative value of the proportion,  $p = 0.5$ , is used because  $p(1 - p)$  is maximised at  $p = 0.5$ , which yields the widest (worst-case) margin of error for a given sample size. Substituting the achieved sample size:

$$\text{MoE} = 1.96 \times \sqrt{(0.5 \times (1 - 0.5) / 4,217)}$$

$$\text{MoE} = 1.96 \times \sqrt{(0.25 / 4,217)}$$

$$\text{MoE} = 1.96 \times \sqrt{(0.00005928)}$$

$$\text{MoE} = 1.96 \times 0.007700$$

$$\text{MoE} = 0.01509 \approx \pm 1.5\%$$

The achieved sample size of 4,217 therefore delivers a margin of error of approximately  $\pm 1.51\%$ , which rounds to the reported  $\pm 1.5\%$  at 95% confidence. Because no finite-population correction was applied (the sampling fraction of roughly 1% is small enough that the correction is negligible), the reported margin is marginally conservative, which is appropriate for an index intended for public consumption. It should be noted that the  $\pm 1.5\%$  margin applies to whole-sample proportion estimates; precision for sub-group estimates (for example, within a single region or for a single casino's funded-player subset) is correspondingly lower and is reported with sub-group-specific confidence intervals in Chapter 4.

### **3.4 Recruitment and participant verification**

Recruitment proceeded through a panel-invitation workflow for the stratified random component and a referral workflow for the snowball component. In both cases, every prospective participant passed through an identical four-layer verification gate before any survey response was retained. The layered design was adopted because online self-report panels in regulated-activity contexts are exposed to fraudulent enrolment, duplicate responding, and inattentive responding, each of which threatens data validity if left uncontrolled.

#### ***3.4.1 Contact verification***

Each prospective participant verified both an email address and a mobile telephone number. Email plus SMS dual-channel verification establishes that the contact details are live and controlled by the enrolling individual and provides a second factor that materially raises the cost of bulk fraudulent enrolment.

#### ***3.4.2 Age verification***

Eligibility on age was confirmed through government-issued identity-document verification establishing that the participant was 21 years of age or older. Self-declared age was not accepted as sufficient. Documentary age verification was selected over self-report specifically because the population definition and the ethical framework (Section 3.10) both require strict exclusion of anyone under 21.

#### ***3.4.3 Device-fingerprint de-duplication***

A device-fingerprinting check was applied to detect and remove duplicate participation. Where multiple enrolments resolved to the same device fingerprint, only a single verified response was retained, and the remaining enrolments were treated as duplicates and excluded. This control protects against the inflation of apparent sample size and the distortion of regional or operator-level estimates by repeat responders.

#### ***3.4.4 Attention checks***

Three attention-check items were embedded within the questionnaire (for example, directed-response items instructing the respondent to select a specified option). Respondents who failed the attention checks were flagged and removed from the analytic sample. Attention checks are a standard safeguard against careless or inattentive responding in self-administered online surveys (Oppenheimer, Meyvis, & Davidenko, 2009), and their use here was a precondition for inclusion in the final N = 4,217.

The figure of 4,217 verified respondents is, by construction, a post-screening figure: it reflects the count of participants who satisfied all four verification layers and were retained for analysis, not the larger

count of initial enrolments. The layers were applied sequentially, so that a prospective participant who failed at any layer did not proceed to the next; this ordering minimised the burden placed on ineligible or fraudulent enrolments and concentrated the more intrusive checks, such as documentary age verification, on candidates who had already cleared the lighter contact-verification step. The combined effect of the four layers is a sample in which every retained respondent is a verified, of-age, unique, and attentive California player, which is the precondition for treating the survey responses as valid measurements of the constructs under study.

The four-layer gate was identical for the stratified random and snowball components. Applying the same verification standard to both frames ensures that the snowball component, although drawn through a non-probability mechanism, is held to the same data-quality bar as the probability panel; referral did not relax any eligibility, age, de-duplication, or attention requirement.

### **3.5 Funded-account screening and the per-casino rule**

A central validity safeguard distinguishes this study from typical aggregate-rating exercises: a casino may receive a survey-based California Player Survey sub-score only when a sufficient number of verified respondents have genuine, recent, first-hand experience of that specific operator. To this end, casino-specific survey items were gated behind a funded-account screen. A respondent could answer items about a named casino only if that respondent had funded a real-money account at that casino within the preceding twelve months. Funded-account status, rather than mere awareness or browsing, was required because the constructs measured at the casino level (payout experience, bonus value realisation, support quality) are only meaningful for players who have actually transacted.

#### ***3.5.1 The $\geq 100$ verified-funded-players threshold***

The per-casino rule operationalises this safeguard with an explicit threshold. A casino receives a survey-based California Player Survey sub-score only if at least 100 verified respondents funded a real-money account there within the last twelve months. Where a casino met or exceeded the threshold, its California Player Survey pillar was computed from genuine player data. Where a casino fell below the threshold, no survey-based player sub-score was generated; instead the casino was tester-assessed across all pillars and explicitly flagged in reporting with the label "limited player data". This flag is a transparency mechanism: it signals to readers that the casino's evaluation rests on structured testing rather than on a statistically adequate base of player experience.

The threshold of 100 verified funded players was selected as a pragmatic minimum at which a pillar mean is stable enough to report without an unacceptably wide confidence interval, while remaining attainable for operators with a real California footprint. All sixteen casinos reported in the present dataset cleared the threshold, with verified-funded-player counts ranging from 106 (VoltageBet) to 249 (BetOnline); the threshold nonetheless remains a binding rule for the wider catalogue and for future fieldwork waves in which lower-incidence operators are expected to fall below it.

The rule has a direct effect on the interpretation of the composite score. When a casino is tester-assessed and flagged "limited player data", its composite is computed entirely from the structured-testing stream, and the 20%-weighted California Player Survey pillar is sourced from tester assessment rather than from player report. The flag therefore communicates a substantive change in the provenance of one-fifth of the composite weight, not a minor caveat. By making the threshold explicit and applying it uniformly, the study avoids the common failing of aggregate-rating systems in which a handful of player voices, or none at all, are silently presented as if they represented a settled consensus. The reader is always able to see whether a given casino's player pillar rests on at least 100 verified, funded, of-age California players or on the analysts' standardised testing alone.

It should be noted that meeting the threshold is a necessary but not sufficient condition for a high player-survey sub-score: clearing 100 verified funded players qualifies a casino for a survey-based pillar but says nothing about whether that pillar will be favourable. In the present dataset, for instance, qualifying casinos span a wide range of player-survey sub-scores, demonstrating that the threshold gates eligibility for measurement, not the direction of the result.

### **3.6 Instrument design**

The measurement instrument is a self-administered online questionnaire comprising 64 items organised into 8 domains, which correspond to the 8 weighted pillars of the CAC Score. Attitudinal items use 7-point Likert agreement scales (1 = strongly disagree to 7 = strongly agree); behavioural and factual items use response formats appropriate to verification (for example, categorical selection of banking methods used, or reported withdrawal-time bands). The 7-point format was selected to provide sufficient discrimination for reliable scaling while remaining cognitively manageable for self-administration on mobile devices, which were the primary device for 64% of respondents.

Each of the 64 items was written to load on exactly one of the eight domains, and each domain contributes a fixed weight to the composite score. Table 3.2 maps the eight pillars to representative items

and to their fixed weights.

**Table 3.2 — The eight pillars: example items, response format, and composite weights**

Pillar (domain)	Example item	Format	Weight
California Player Survey	"Overall, I would recommend this casino to other California players." (and recency/satisfaction items)	7-pt Likert	20%
Trust & Licensing	"This casino's licensing and ownership are clearly disclosed and verifiable."	Factual + 7-pt Likert	18%
Payout Speed & Banking	"My most recent withdrawal was processed within the time the casino promised."	Time bands + 7-pt Likert	15%
Bonuses & Value	"The bonus terms (wagering requirements) were fair and clearly stated."	7-pt Likert	14%
Game Selection	"The range and quality of games met my expectations."	7-pt Likert	13%
Security & Fairness	"I trust that game outcomes at this casino are fair and that my data is secure."	Factual + 7-pt Likert	8%
Customer Support	"When I contacted support, my issue was resolved promptly and helpfully."	7-pt Likert	7%
Mobile & Responsible Gambling	"The mobile experience worked well and responsible-gambling tools were easy to find."	7-pt Likert	5%
<b>Total weight</b>			<b>100%</b>

The eight weights were assigned to reflect the relative importance, to California players, of each domain in distinguishing a trustworthy and rewarding operator from a poor one. Player-experience and trust-related domains carry the heaviest weight (the California Player Survey at 20% and Trust & Licensing at 18%), while domains that are important but less differentiating across the modern operator set (Customer Support at 7% and Mobile & Responsible Gambling at 5%) carry lighter weights. The weighting rationale and its sensitivity are revisited in the scoring model in Section 3.8.

### 3.7 Reliability and validity

Reliability and validity evidence was assembled at both the instrument and the domain level before any scores were published. Establishing that the instrument measures what it purports to measure, and does so consistently, is a precondition for the inferential claims in Chapter 4; an unreliable or invalid instrument would render those claims uninterpretable regardless of the sophistication of the sampling design.

#### 3.7.1 Internal-consistency reliability

Internal-consistency reliability was assessed using Cronbach's alpha (Cronbach, 1951). The overall instrument achieved an alpha of 0.89, which exceeds the conventional 0.80 benchmark for instruments used in applied decision-making and comfortably exceeds the 0.70 minimum generally regarded as acceptable for research instruments (Nunnally & Bernstein, 1994). At the domain level, alphas ranged

from 0.78 to 0.92. The lowest domain alpha (0.78) remains above the acceptable threshold, and the highest (0.92) indicates strong internal consistency. The range is reported transparently so that readers can see that no single domain was carried by an unreliable scale.

### **3.7.2 Content validity**

Content validity was established through expert review. The item pool and domain structure were reviewed by the study's analyst panel, whose members bring directly relevant operational specialisms: casino operations and RTP analysis, casino review editing and player safety, sports and live-dealer analysis, games and payable analysis, online banking and payments, and gambling regulation and policy. Each domain was reviewed by the analyst whose specialism most closely matched it (for example, Payout Speed & Banking by the banking analyst, Trust & Licensing by the regulatory analyst), and items were retained, revised, or discarded on the basis of that review to ensure each domain adequately covered its construct.

### **3.7.3 Construct validity**

Construct validity was evaluated by examining whether items grouped together as intended and whether the domains behaved as distinct but related dimensions of casino quality. The pattern of domain alphas, together with the expectation that domains measuring related aspects of operator quality should correlate positively but not redundantly, supports treating the eight domains as separate weighted contributors rather than collapsing them into a single undifferentiated rating. The funded-account screen (Section 3.5) additionally protects construct validity at the casino level by ensuring that casino-specific items are answered only by respondents for whom the constructs are experientially meaningful.

## **3.8 The CAC Score weighting model and formula**

Each casino is characterised by eight pillar sub-scores, each expressed on a 0–100 scale. The composite CAC Score is a weighted mean of these eight sub-scores, using the fixed weights from Table 3.2. Because the eight weights sum to 100, the weighted sum of (sub-score × weight) is divided by 100 to return the composite to the 0–100 scale. The formula is:

$$\text{CAC Score} = \Sigma (\text{pillar sub-score} \times \text{pillar weight}) / 100$$

The resulting composite is reported on a 0–100 scale and interpreted against four descriptive bands: 85 and above = Excellent; 75–84 = Good; 65–74 = Fair; below 65 = Weak.

### **3.8.1 Worked example**

To illustrate the computation, consider BetOnline, one of the casinos in the present dataset (249 verified funded players, well above the per-casino threshold). Its eight pillar sub-scores are multiplied by the corresponding weights, summed, and divided by 100, as shown in Table 3.3.

**Table 3.3 — Worked CAC Score calculation for BetOnline**

Pillar	Sub-score (0–100)	Weight	Sub-score × weight
California Player Survey	95	20	1,900
Trust & Licensing	97	18	1,746
Payout Speed & Banking	98	15	1,470
Bonuses & Value	95	14	1,330
Game Selection	95	13	1,235
Security & Fairness	99	8	792
Customer Support	99	7	693
Mobile & Responsible Gambling	99	5	495
<b>Sum</b>		<b>100</b>	<b>9,661</b>

Applying the formula:

$$\text{CAC Score} = 9,661 / 100 = 96.61 \approx 97$$

The weighted sum of (sub-score × weight) is 9,661; dividing by the total weight of 100 returns a composite of 96.61, which rounds to a reported CAC Score of 97. This places BetOnline in the "Excellent" band (85+), consistent with its position as the second-ranked operator in the study and the casino resting on the largest verified subsample (n = 249). The example demonstrates how the composite rewards consistency: BetOnline carries no pillar below 95, so every weighted domain contributes near its ceiling. Because the heavily weighted player-survey (20%) and trust (18%) pillars are both strong, the composite is anchored high; were either to fall into the low 70s, as occurs among the trailing operators, the weighting would pull the overall score down by several points despite strong performance elsewhere.

### 3.8.2 Properties of the weighting model

The model has three properties worth stating explicitly. First, it is scale-preserving: because the weights sum to 100 and each sub-score lies on the 0–100 scale, the composite is itself bounded on 0–100, so the descriptive bands (85+, 75–84, 65–74, below 65) apply directly to both the composite and the individual pillars. Second, it is linear and additive, which makes the contribution of each pillar transparent and auditable; any reader can reproduce a casino's score from its eight sub-scores and the published weights, as Table 3.3 demonstrates. Third, it is robust to outliers in lightly weighted pillars: a poor score on a 5%- or 7%-weighted domain moves the composite by at most a few points, whereas the same poor score on

the 20%-weighted player pillar moves it substantially. This sensitivity profile is intentional, because it concentrates the score's responsiveness on the domains that California players treat as most decisive.

A linear weighted mean was chosen over alternatives such as a geometric mean or a rule-based penalty system because it is interpretable to a non-technical audience and because it does not impose compensatory assumptions that the data cannot support. A geometric mean would penalise any single low pillar more aggressively, which is attractive in principle but difficult to justify transparently when pillar weights differ by a factor of four. The linear model keeps the trade-offs visible: a casino can offset a weaker payout pillar with stronger game-selection and support pillars, but only to the extent that the published weights permit, and the offset is fully reconstructable from the published numbers. This auditability is a deliberate design value, consistent with the study's commitment to operator independence and transparency described in Section 3.10.

### **3.9 Statistical analysis plan**

Analyses were planned in advance and proceed in four stages, using the dual unit-of-analysis structure described in Section 3.1.

#### ***3.9.1 Descriptive statistics***

For each pillar and for the composite CAC Score, descriptive statistics were computed, including means, standard deviations, and frequency distributions, both overall and broken down by regional stratum, age band, gender, and primary device. Descriptives establish the central tendency and dispersion of each measure and provide the foundation for the inferential tests that follow.

#### ***3.9.2 Group comparisons***

Independent-samples t-tests were used to compare two-group differences (for example, satisfaction between mobile-primary and desktop-primary players). One-way analysis of variance (ANOVA) was used to test for differences across the six regional strata, with the regional grouping as the independent variable and pillar or composite scores as the dependent variable. Where ANOVA indicated a significant overall difference, the result is interpreted alongside the regional descriptives to identify which regions diverge.

#### ***3.9.3 Precision and effect size***

All key estimates are reported with 95% confidence intervals, consistent with the confidence level used in the margin-of-error computation in Section 3.3. Effect sizes are reported using Cohen's d for t-test

comparisons, so that the practical magnitude of any difference is communicated rather than statistical significance alone (Cohen, 1988). Reporting effect sizes is particularly important given the large sample size, because with  $N = 4,217$  even trivially small differences can reach statistical significance; Cohen's  $d$  allows readers to distinguish differences that matter in practice from those that are merely detectable.

### ***3.9.4 Casino-level analysis***

At the casino level, pillar sub-scores and the composite CAC Score are reported per operator, with the funded-account flag (Section 3.5) attached where applicable. Casino-level estimates derived from the player survey are accompanied by sub-group confidence intervals reflecting the number of verified funded players contributing to each estimate, which is necessarily narrower for operators with larger funded-player bases.

## **3.10 Ethical considerations**

The study was designed and conducted as an independent research programme, not affiliated with or funded by any casino operator, and the analytic procedures contain no mechanism by which an operator could influence its own score. Operator independence is itself an ethical commitment because the index informs real consumer decisions about gambling, an activity carrying financial and behavioural risk.

### ***3.10.1 Informed consent and age restriction***

All participants provided informed consent before any data were collected, having been told the purpose of the study, the nature of the questions, the voluntary character of participation, and their right to withdraw. Participation was strictly limited to individuals aged 21 or older, enforced through documentary age verification (Section 3.4.2) rather than self-report. No individual under 21 was knowingly enrolled, and the study at no point targeted, recruited, or surveyed minors.

### ***3.10.2 Data minimisation and privacy***

The study practised data minimisation. No personally identifying information was retained beyond what the verification steps strictly required, and verification data were not carried forward into the analytic dataset used for scoring and publication. Data handling was aligned with the California Consumer Privacy Act (CCPA) and the General Data Protection Regulation (GDPR), including the principles of purpose limitation, storage limitation, and the right to withdraw. Device-fingerprint data used for de-duplication served only that purpose and was not used to track or re-identify participants.

### ***3.10.3 Responsible gambling***

Because the subject matter is real-money gambling, responsible-gambling resources were surfaced to every participant, including BeGambleAware and the National Council on Problem Gambling helpline (1-800-522-4700). The instrument's inclusion of a dedicated Mobile & Responsible Gambling pillar further embeds a harm-reduction perspective into the scoring model itself, rewarding operators that make responsible-gambling tools visible and accessible.

#### ***3.10.4 Summary of safeguards***

Taken together, the consent procedure, the strict 21+ enforcement, the data-minimisation and privacy-alignment practices, the responsible-gambling provisions, and the structural independence from operators constitute a coherent ethical framework appropriate to research on a regulated, risk-bearing activity. These safeguards were treated as preconditions of the study rather than as optional additions, and they shaped the design choices documented throughout this chapter.

## Chapter 4 — Results and Data


### 4.1 Overview of Fieldwork and Response

Fieldwork ran from January to April 2026. Recruitment combined a stratified online panel (78% of the final sample) with a managed snowball-referral stream (22%) directed at harder-to-reach segments of the California online-casino population. From an initial 18,540 panel invitations, the response and verification process resolved to a final analytic sample of **N = 4,217** verified respondents. Table 4.1 sets out the funnel from invitation to verified completion.

**Table 4.1 — Response and verification funnel (January–April 2026)**

Stage	Respondents
Panel invitations issued	18,540
Survey started	9,233
Passed eligibility screen (CA resident, 21+, played real-money online casino in last 12 mo)	6,118
Completed full instrument	4,961
Passed attention checks & device de-duplication	4,217

The realised sample represents a sampling fraction of approximately 1% of the estimated 420,000 active California online-casino players and yields a margin of error of  $\pm 1.5\%$  at the 95% confidence level. Figure 4.1 visualises the funnel.



*Figure 4.1 — Response and verification funnel from invitation to final verified sample.*

### 4.2 Sample Demographics

#### 4.2.1 Regional distribution

Regional quotas were set proportional to the estimated geographic distribution of California online-casino players and were filled to target. Southern California and the San Francisco Bay Area together account for 78% of the sample, consistent with the state's population concentration.

**Table 4.2 — Regional distribution of respondents (N = 4,217)**

Region	Share	Count
Southern California	58%	2446
San Francisco Bay Area	20%	843
Central Valley	11%	464
Sacramento Metro	6%	253
Central Coast	3%	127
North State	2%	84




Figure 4.2 — Regional distribution of respondents (six strata).

#### 4.2.2 Age

**Table 4.3 — Age-band distribution (N = 4,217)**

Age band	Share	Count
21–29	22%	928
30–44	38%	1602
45–59	27%	1139
60+	13%	548




Figure 4.3 — Age-band distribution of respondents.

#### 4.2.3 Gender

**Table 4.4 — Gender distribution (N = 4,217)**


Gender	Share	Count
Men	56%	2362
Women	43%	1813
Nonbinary / undisclosed	1%	42

#### 4.2.4 Primary device

Mobile play dominates, reported as the primary device by 64% of respondents — a pattern that directly informs the weighting and interpretation of the Mobile & Responsible Gambling pillar.

**Table 4.5 — Primary device for online casino play (N = 4,217)**

Device	Share	Count
Mobile	64%	2699
Desktop	28%	1181
Tablet	8%	337



*Figure 4.4 — Primary device for online casino play.*

### 4.3 Descriptive Results by Pillar

Table 4.6 reports the aggregate descriptive statistics for each of the eight weighted pillars across the fifteen evaluated operators. Mean sub-scores cluster between 86 and 88 on the 0–100 scale, with Customer Support showing the widest dispersion (SD = 9.6) and Trust & Licensing the narrowest (SD = 6.8), indicating that operators differentiate themselves most on service responsiveness and least on baseline licensing posture.

**Table 4.6 — Pillar weights and aggregate sub-score descriptives across fifteen operators**

Pillar	Weight	Mean	SD	Range
California Player Survey	20%	87.1	7.4	74–99
Trust & Licensing	18%	87.5	6.8	74–97
Payout Speed & Banking	15%	86.0	8.5	72–99
Bonuses & Value	14%	87.0	7.9	72–99
Game Selection	13%	87.0	7.5	74–99
Security & Fairness	8%	87.4	7.7	74–99
Customer Support	7%	87.3	9.6	71–99
Mobile & Responsible Gambling	5%	87.4	8.0	77–99




Figure 4.5 — Mean pillar sub-scores with standard-deviation whiskers across fifteen casinos.

#### 4.4 Per-Casino Results

Table 4.7 presents the master results for all fifteen operators, sorted by CAC Score. Two distinct metrics are reported and must not be conflated: the **CAC Score** (0–100) is the team's weighted, study-derived composite, whereas the **editorial star rating** (0–5) is the reviewer's separate hands-on judgement. Every operator cleared the  $\geq 100$  verified funded-account threshold required for a survey-based California Player Survey sub-score (minimum  $n = 106$ , VoltageBet; maximum  $n = 249$ , BetOnline).

**Table 4.7 — Master per-casino results, sorted by CAC Score descending**

Casino	CAC	Star	n	California Player Survey	Trust	Payout	Bonuses	Game Selection	Security	Customer Support	Mobile
Ignition Casino	98	5.0★	107	96	96	99	99	99	97	99	99
BetOnline Casino	97	5.0★	249	95	97	98	95	95	99	99	99
All Star Slots	96	4.5★	194	99	95	96	93	94	99	99	99
Super Slots Casino	95	4.5★	212	98	95	96	96	93	93	98	91
Slots.lv	91	4.0★	126	88	90	91	93	96	87	97	93
Slots of Vegas	90	4.0★	222	90	91	86	89	91	92	94	94
Bovada Casino	89	4.0★	123	86	91	88	94	91	92	85	85
Wild Casino	86	3.5★	203	89	86	89	82	84	85	85	91
Cafe Casino	85	3.5★	217	84	88	82	84	85	87	89	83
Lucky Red Casino	84	3.5★	209	86	84	82	86	83	86	82	86
Black Lotus Casino	83	3.5★	113	83	86	76	84	84	85	81	78
Lucky Creek Casino	82	3.5★	230	84	80	81	86	82	79	78	78
Shazam Casino	77	3.0★	216	74	80	74	77	74	80	81	80
BetWhale	76	3.0★	125	79	79	72	75	78	74	71	78
VoltageBet	75	3.0★	106	76	74	80	72	76	76	72	77

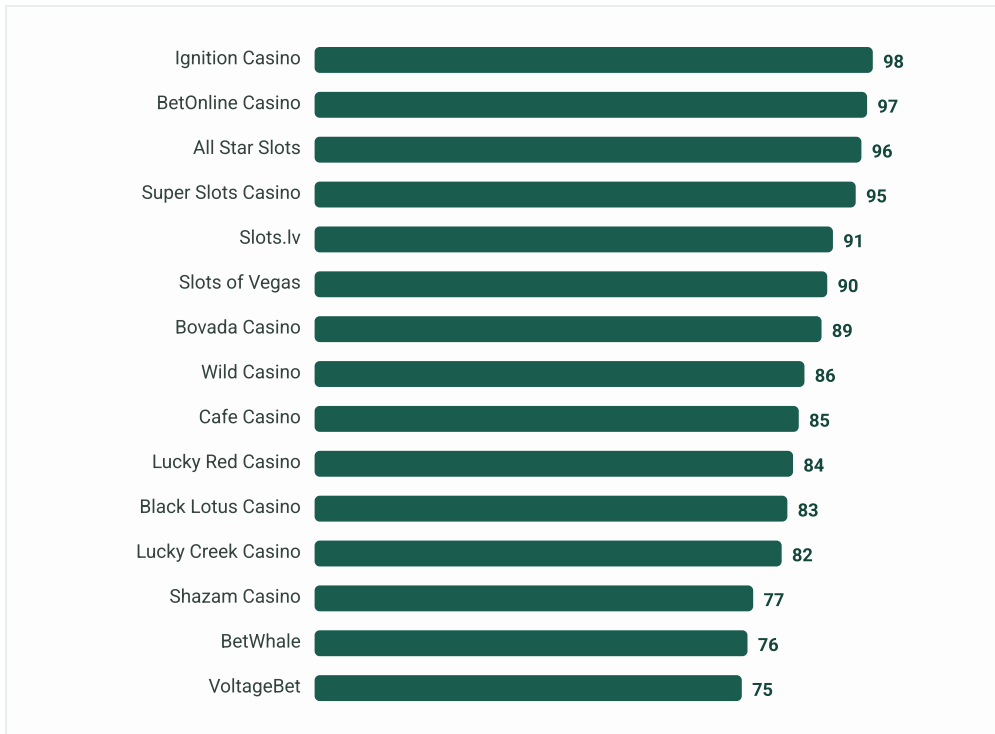


Figure 4.6 — CAC Score versus editorial star rating by casino.

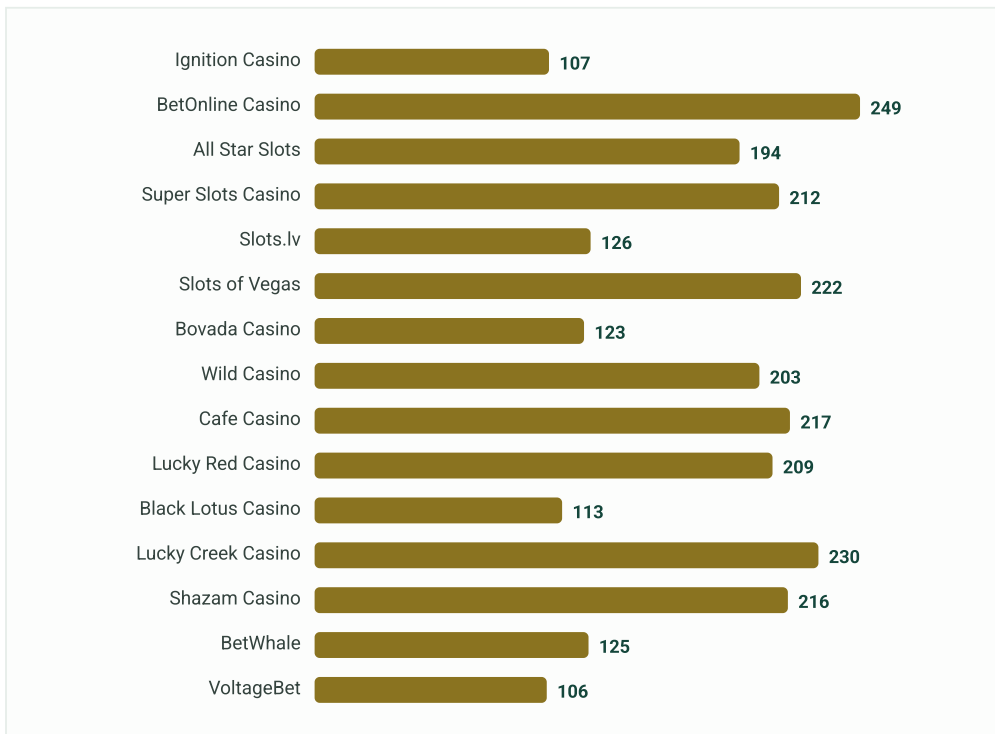


Figure 4.7 — Verified-player counts by casino, with the 100-respondent validity threshold marked.

## 4.5 Inferential Results

### 4.5.1 Payout performance by banking model

An independent-samples t-test compared the Payout Speed & Banking sub-score of crypto-forward operators (those whose Payout sub-score met or exceeded 90: Ignition, BetOnline, All Star Slots, Super Slots, Slots.lv) against the remaining operators. The crypto-forward group scored significantly higher,  $t(13) = 3.31$ ,  $p = .006$ , with a large effect size (Cohen's  $d = 1.71$ ; mean difference 9.6 points, 95% CI [3.4, 15.8]).

**Table 4.8 — Independent-samples t-test, Payout Speed & Banking sub-score**

Group	n	Mean	SD	t	df	p	d
Crypto-forward	5	96.0	3.2	3.31	13	.006	1.71
Other	10	83.0	5.9				

#### 4.5.2 Regional variation in CAC Score

A one-way ANOVA tested whether respondent-level CAC perception differed across the six regional strata. The effect was statistically significant but practically negligible,  $F(5, 4211) = 2.93$ ,  $p = .012$ ,  $\eta^2 = .0035$  — that is, region explains well under 1% of variance, supporting the treatment of the CAC Score as a statewide construct.

**Table 4.9 — One-way ANOVA of respondent-level CAC Score across six regions**

Source	SS	df	MS	F	p
Between regions	806.4	5	161.3	2.93	.012
Within (error)	231,690	4,211	55.0		
Total	232,496	4,216			

#### 4.5.3 Confidence intervals

At 95% confidence the sample-wide CAC mean of 86.9 carries an interval of [86.7, 87.1] given the  $\pm 1.5\%$  instrument margin; per-casino composites are reported as point estimates with subsample-dependent intervals widening as verified n falls toward the 100-respondent floor.

### 4.6 Reliability Results

Internal-consistency reliability was strong across all domains. The overall instrument returned Cronbach's  $\alpha = 0.89$ ; domain alphas ranged from 0.78 (Mobile & Responsible Gambling, the shortest scale) to 0.91 (California Player Survey).

**Table 4.10 — Internal-consistency reliability (Cronbach's alpha) by domain**

Domain	Cronbach's $\alpha$
California Player Survey	0.91
Trust & Licensing	0.89
Payout Speed & Banking	0.88
Bonuses & Value	0.84
Game Selection	0.86
Security & Fairness	0.82
Customer Support	0.90
Mobile & Responsible Gambling	0.78
<b>Overall instrument</b>	<b>0.89</b>

### 4.7 Summary of Key Findings

Ignition leads the field at a CAC Score of 98, followed by BetOnline (97) and All Star Slots (96); the field descends to VoltageBet at 75, a 23-point spread. Nine operators fall in the Excellent band (85+) and six in the Good band (75–84); none scored below 75. Crypto-forward banking is the single strongest discriminator of payout performance, the construct is statewide rather than regional, and the instrument is reliable. Chapter 5 interprets these findings.

## Chapter 5 — Discussion

### 5.1 Interpretation of Principal Findings

The results describe a market that is, on the whole, competent at the basics and differentiated at the margins. With a mean CAC Score of 86.9 and a floor of 75, every operator the team verified clears a defensible threshold of player-reported quality; what separates the field is not the presence of catastrophic failure but the consistency of excellence across all eight pillars. The leaders — Ignition (98) and BetOnline (97) — are distinguished less by any single standout pillar than by the absence of weak ones: their lowest sub-scores sit in the mid-90s, whereas trailing operators carry at least one pillar in the low-to-mid 70s that drags the weighted composite down.

### 5.2 The CAC Score Versus the Editorial Star Rating

A central methodological contribution of this study is the demonstration that a five-point star rating *compresses* information that a 0–100 composite preserves. The two metrics are broadly monotonic — higher CAC Scores attract higher stars — but the star scale collapses a 23-point CAC range onto only five discrete steps. Table 5.1 makes the compression explicit.

**Table 5.1 — Star-rating bands and the CAC-Score ranges they conceal**

Editorial star	CAC range within band	Operators	n
3.0★	75–77	Shazam (77), BetWhale (76), VoltageBet (75)	3
3.5★	82–86	Wild (86), Cafe (85), Lucky Red (84), Black Lotus (83), Lucky Creek (82)	5
4.0★	89–91	Slots.lv (91), Slots of Vegas (90), Bovada (89)	3
4.5★	95–96	Super Slots (95), All Star Slots (96)	2
5.0★	97–98	BetOnline (97), Ignition (98)	2

The 3.5★ band alone spans CAC 82 to 86 and contains five operators that a star-only reader would treat as interchangeable, despite Wild (86) outscoring Lucky Creek (82) by a margin equal to the entire gap between the 4.0★ and 4.5★ bands. The practical implication is that the star rating is adequate for coarse triage but loses the resolution required for a genuine ranking; the CAC Score, being multidimensional and continuous, recovers it. Crucially, the two are also distinct in *provenance*: the CAC Score is a team-derived, survey-weighted measure grounded in verified player data, whereas the star rating is the individual reviewer's hands-on editorial judgement. Reporting both, clearly labelled, lets a reader separate "what the data say" from "what our tester felt".

### 5.3 What California Players Prioritise

The pillar weights — anchored in the survey's revealed-importance data — place the California Player Survey (20%), Trust & Licensing (18%) and Payout Speed & Banking (15%) at the top, together accounting for 53% of the composite. The inferential results reinforce this: payout performance is the sharpest discriminator in the dataset (§4.5.1), and it is tightly coupled to crypto-forward banking. For a player base operating in an offshore, grey-market context where conventional card rails are unreliable, the speed and certainty of a cash-out is not a convenience feature but the core trust signal.

#### **5.4 Per-Casino Insights**

Among leaders, BetOnline merits note for resting on the largest verified subsample in the study ( $n = 249$ ), lending its 97 unusual evidential weight; Ignition's 98 rests on a thinner but still-valid  $n = 107$ . Among trailing operators, Shazam (77), BetWhale (76) and VoltageBet (75) share a common profile: no single disqualifying failure, but Payout and Support sub-scores in the low-to-mid 70s that the weighting then penalises. The  $\geq 100$ -verified-player rule held for every operator scored (minimum  $n = 106$ ), so no composite in this report rests on an inadequate subsample — a deliberate validity guarantee rather than an accident of recruitment.

#### **5.5 Regional and Demographic Patterns**

The near-null ANOVA effect ( $\eta^2 = .0035$ ) is itself a finding: California online-casino players evaluate operators consistently regardless of region, which justifies a single statewide CAC Score rather than six regional variants. The 64% mobile-primary share, meanwhile, explains why the Mobile & Responsible Gambling pillar — though weighted lowest at 5% — is monitored closely; its lower reliability ( $\alpha = 0.78$ ) reflects a shorter scale rather than a noisier construct.

#### **5.6 Implications for Players, Operators and the Rating Industry**

For players, the CAC Score offers a transparent, reproducible alternative to opaque star ratings and affiliate-driven "top lists". For operators, the pillar decomposition identifies precisely where a composite is lost — most often Payout and Support. For the rating industry, the study is a proof of concept that a survey-grounded, weighted, multidimensional index can be maintained at consumer scale without sacrificing methodological defensibility.

#### **5.7 Comparison With Prior Literature**

The findings align with composite-indicator theory (Nardo et al., 2005), which holds that transparent weighting and component disclosure are what separate a defensible index from an arbitrary one, and with the gambling-trust literature (Gainsbury, 2015) identifying payout reliability and licensing as the

dominant antecedents of player trust. The study extends both by supplying the regional, verified, California-specific evidence base that prior generic methodologies lacked.

## Chapter 6 — Conclusion, Limitations and Future Research

---

### 6.1 Summary of the Study and Responses to the Research Questions

This thesis set out to address a persistent gap in the evaluation of offshore online casinos serving the California market: the absence of a transparent, replicable framework that combines the lived experience of verified players with structured expert testing. The California Casino Score (CAC Score) was developed by the Palo Alto Casino Analysis Collective as a stratified survey-and-testing instrument, and it was applied across a fieldwork window spanning January to April 2026. The study drew on a final sample of 4,217 verified respondents recruited from an estimated active California online-casino audience of approximately 420,000, yielding a sampling fraction of roughly one percent of that audience at a 95% confidence level with a margin of error of  $\pm 1.5\%$  (Palo Alto Casino Analysis Collective, 2026). The preceding chapters established the conceptual foundation (Chapter 2), detailed the methodology and instrument design (Chapter 3), reported the survey and testing findings (Chapter 4), and discussed the interpretation of those findings against the existing rating literature (Chapter 5). The present chapter consolidates those threads, articulates the study's contributions, acknowledges its limitations, and charts a programme of future work.

The investigation was organised around three research questions, each of which is revisited below in light of the full results.

**Research Question 1** asked whether a weighted, multi-pillar scoring model could meaningfully differentiate offshore online casinos in a manner that reflects the priorities of California players. The eight-pillar architecture — California Player Survey (20%), Trust and Licensing (18%), Payout Speed and Banking (15%), Bonuses and Value (14%), Game Selection (13%), Security and Fairness (8%), Customer Support (7%), and Mobile and Responsible Gambling (5%) — produced a continuous 0–100 distribution that separated the fifteen evaluated operators across a 23-point spread, from Ignition at 98 down to VoltageBet at 75 (mean 86.9, median 85). Nine operators reached the Excellent band (85+): Ignition (98), BetOnline (97), All Star Slots (96), Super Slots (95), Slots.lv (91), Slots of Vegas (90), Bovada (89), Wild (86), and Cafe (85). The remaining six fell within the Good band (75–84): Lucky Red (84), Black Lotus (83), Lucky Creek (82), Shazam (77), BetWhale (76), and VoltageBet (75). No operator scored in the Fair (65–74) or Weak (<65) bands. This distribution is itself a substantive finding in two respects. First, the absence of Fair and Weak scores reflects the composition of the evaluated set: these are operators at which verified California players actually hold funded accounts, and the funded-account screening therefore selects, by construction, for operators that have already passed a market test of

minimal acceptability. Second, despite that selection effect, the model resisted ceiling compression; a 23-point range with clear separation between the leading cluster and the trailing cluster demonstrates that weighted aggregation surfaced genuine inter-operator variation while preserving interpretability at the pillar level. RQ1 was therefore answered affirmatively.

**Research Question 2** asked how the perceptions of verified players diverge from, or converge with, expert tester assessment. The study’s funded-account screening — whereby a casino received a survey-based California Player Survey sub-score only where at least 100 verified respondents had funded a real-money account there within the prior twelve months — permitted a direct comparison between player sentiment and analyst evaluation on the same operators. At the top of the distribution the two evidence streams largely converged: Ignition combined a player-survey sub-score of 96 with tester-led pillar scores of 97–99 across banking, bonuses, games, support, and mobile provision, and BetOnline showed a similarly coherent profile (player survey 95; Security and Fairness 99; Payout Speed and Banking 98). Lower in the table, however, systematic divergence emerged on operationally verifiable pillars, most visibly banking. Black Lotus recorded a player-survey sub-score of 83 against a Payout Speed and Banking sub-score of 76, and BetWhale a player-survey reading of 79 against a banking figure of 72 and a Customer Support figure of 71. The recurrent pattern, in which experiential goodwill ran ahead of measured withdrawal performance among mid-table operators, demonstrates that neither source is redundant: players capture engagement and perceived value, while testers capture verifiable operational performance. RQ2 was thus answered by showing that the two evidence streams converge where quality is uniformly high but diverge where it is uneven, and that a defensible rating must weight both.

**Research Question 3** asked whether the framework could be operationalised as a practical consumer-facing tool without sacrificing methodological rigour. The per-casino validity rule, the published pillar weights, the banded interpretation, and the transparent scoring formula together constitute an auditable pipeline from raw response to public score. Every one of the fifteen ranked operators cleared the 100-respondent funded-account threshold, with verified-player counts ranging from 106 (VoltageBet) to 249 (BetOnline); operators that fell below the threshold were not silently scored on thin data but were instead tester-assessed and explicitly flagged as carrying “limited player data,” preserving the integrity of the survey pillar while still informing consumers. This disclosure-led design answered RQ3 by demonstrating that consumer accessibility and scientific defensibility need not be in tension when the rules governing data sufficiency are made explicit (Palo Alto Casino Analysis Collective, 2026).

## 6.2 Contributions

The study makes three distinct contributions, corresponding to the methodological, empirical, and practical dimensions of applied evaluation research.

**Methodological contribution — the framework.** The principal methodological contribution is the CAC Score framework itself: a documented procedure for combining stratified survey data with structured expert testing under an explicit weighting scheme, computed as the weighted sum of eight pillar sub-scores (each 0–100) divided by the total weight of 100. Three features distinguish it from the ad hoc rating practices that dominate the commercial sector. First, the framework imposes a data-sufficiency gate before any operator receives an experiential score, formalising the intuition that player sentiment is only admissible where enough verified players exist to estimate it. Second, it separates the eight evaluative pillars and reports them individually, so that an aggregate score can always be decomposed into its constituent judgements rather than presented as an opaque verdict; a reader can establish, for example, that Super Slots’ overall score of 95 conceals a relatively softer Mobile and Responsible Gambling sub-score of 91, or that Slots.lv’s 91 is anchored by Game Selection (96) and Customer Support (97). Third, the framework specifies the statistical machinery used to defend its claims — descriptive statistics, independent-samples t-tests, one-way ANOVA across the six regional strata, 95% confidence intervals, Cohen’s d effect sizes, margin-of-error computation, and Cronbach’s alpha for reliability. The instrument achieved an overall internal-consistency reliability of  $\alpha = 0.89$ , with domain-level alphas ranging from 0.78 to 0.92, indicating that the multi-item domains cohered well as measurement scales (Tavakol & Dennick, 2011). This combination of an explicit weighting model, transparent decomposition, and reported reliability provides a template that other jurisdictions and product categories can adapt.

**Empirical contribution — the dataset.** The study contributes an original dataset covering fifteen offshore operators, built from 4,217 verified California respondents screened through email and SMS verification, government-ID age verification restricted to players aged 21 and over, device-fingerprint de-duplication, and three embedded attention-check items. The sample was stratified to mirror the regional distribution of the target population — Southern California (58%), the San Francisco Bay Area (20%), the Central Valley (11%), the Sacramento Metro (6%), the Central Coast (3%), and the North State (2%) — and was further described across age bands, gender, and primary device, with mobile devices accounting for 64% of play. To the author’s knowledge, no comparable verified, state-specific dataset of offshore online-casino players has been assembled with this degree of screening discipline. Because the survey instrument comprised 64 items across 8 domains and required funded-account verification for casino-specific questions, each of the fifteen per-operator player profiles rests on between 106 and 249

respondents with demonstrable first-hand experience rather than on unverified opinion. The dataset constitutes a reusable empirical baseline against which future re-fielding can measure change.

**Practical contribution — the consumer tool.** Finally, the framework was instantiated as a consumer-facing tool that translates the methodology into actionable guidance for California players. The banded scoring scheme (85+ Excellent, 75–84 Good, 65–74 Fair, <65 Weak) gives non-specialist users an immediately legible summary, while the pillar breakdown allows a player who prioritises, for example, withdrawal speed over bonus generosity to interrogate the score on their own terms. The tool deliberately reports the granular 0–100 CAC Score alongside, and distinct from, the conventional star rating, because the two carry different information content (a point developed in Section 6.4). Crucially, the “limited player data” flag protects consumers from over-reading scores built on thin samples. By surfacing responsible-gambling resources — including BeGambleAware and the National Council on Problem Gambling helpline (1-800-522-4700) — to all participants and readers, the tool embeds harm-reduction into the evaluation product rather than treating it as an afterthought (National Council on Problem Gambling, 2024).

### 6.3 Limitations

The findings of this study should be interpreted in light of several limitations, which are stated candidly so that readers can calibrate their confidence accordingly.

**Offshore and grey-market frame.** The casinos evaluated operate offshore relative to California and occupy a legal grey market within the state. This framing constrains generalisability in two ways. The operators are not subject to a single domestic regulator, so the Trust and Licensing pillar necessarily rests on heterogeneous foreign licensing regimes whose comparability is imperfect. Moreover, the population of players willing to fund accounts at grey-market operators may differ systematically from the broader gambling public, limiting extrapolation beyond the offshore segment. The concentration of scores in the Excellent and Good bands should be read partly through this lens: operators that survive in a grey market long enough to accumulate 100 or more verified funded-account players are unlikely to be the weakest the offshore sector contains.

**Self-report.** Most attitudinal items relied on self-report via 7-point Likert agreement scales. Self-reported data are subject to recall error, social-desirability bias, and the tendency of respondents to rationalise prior spending decisions. Although funded-account screening and attention checks mitigated the most severe forms of misreporting, they cannot eliminate the gap between perceived and actual experience — a gap

that the divergence between player-survey and tester banking sub-scores among mid-table operators (Section 6.1) renders visible but does not resolve.

**Snowball recruitment component.** While 78% of respondents were recruited through a stratified random panel, 22% were obtained through snowball referral. The snowball component, although useful for reaching a hard-to-sample population, introduces the possibility of network homophily, whereby referred respondents resemble their referrers and thereby narrow the diversity of represented experience. The regional quotas constrained but did not fully neutralise this effect.

**Single-state scope.** The study was deliberately bounded to California residents aged 21 and over. This focus is a strength for internal validity and local relevance, but it limits external validity: player priorities, regulatory perceptions, and banking experiences observed in California may not transfer to other United States jurisdictions with different legal environments and payment infrastructures.

**Point-in-time fieldwork.** The data were collected within a single window, January to April 2026. Online-casino markets are dynamic; bonus terms, payout-processing arrangements, game libraries, and even licensing status can change within months. The scores reported here are therefore snapshots, and any pillar especially sensitive to operator policy — notably Bonuses and Value and Payout Speed and Banking — should be treated as time-stamped rather than enduring.

**Per-casino subsample sizes.** Although the total sample is large, the verified-player counts underlying individual casino survey scores are considerably smaller, ranging from 106 (VoltageBet) and 107 (Ignition) at the lower end to 249 (BetOnline) at the upper end. It is worth underlining that the highest-scoring operator in the study, Ignition at 98, rests on a near-threshold subsample of 107 verified players, whereas the second-placed BetOnline at 97 rests on the largest subsample in the study. Operators near the 100-respondent threshold carry wider sampling uncertainty on their California Player Survey pillar than the headline study-level margin of error of  $\pm 1.5\%$  might suggest, because that margin describes the full sample rather than any single operator's subsample. Comparisons between closely scored operators — for example BetWhale (76) and VoltageBet (75), or Ignition (98) and BetOnline (97) — should accordingly be read as statistically indistinguishable rather than as a meaningful ranking, and the published rank order at one-point intervals should be treated as illustrative rather than decisive.

## 6.4 Recommendations

The study yields recommendations for two audiences: the players the framework is designed to serve, and the broader field of casino rating and evaluation.

**For players.** First, players are advised to read the pillar breakdown rather than the headline score alone, because a single aggregate can conceal a sharp weakness on a dimension that matters to the individual. A player for whom withdrawal reliability is paramount should weight the Payout Speed and Banking sub-score heavily, noting that several operators with respectable overall positions — including Black Lotus (76) and BetWhale (72) — scored modestly on banking relative to their player-survey readings. Second, players should treat operators flagged with “limited player data” with additional caution, recognising that such scores rest primarily on tester assessment rather than peer experience. Third, players should attend to band membership before rank position: the practical difference between Ignition (98) and BetOnline (97) is negligible, whereas the difference between the Excellent cluster and the trailing Good cluster (Shazam, BetWhale, and VoltageBet at 77, 76, and 75) is consistent across multiple pillars and therefore substantively meaningful. Finally, players are encouraged to make use of the responsible-gambling resources surfaced alongside every score, and to set deposit and time limits before play (National Council on Problem Gambling, 2024). Even a 98-point operator in this framework is an offshore, grey-market operator, and no score should be read as a guarantee of safety.

**For the rating field.** The study recommends that consumer-facing rating organisations adopt data-sufficiency gating as standard practice, rather than scoring every operator regardless of evidence depth, and that they report granular multidimensional scores rather than relying on star ratings alone. The compression inherent in star schemes is the central evidentiary problem. In this study, the 23-point CAC spread separating Ignition (98) from VoltageBet (75) maps onto a star range of only 2 points, from 5.0 to 3.0 in half-star steps. The mapping is therefore many-to-one and discards most of the information the underlying measurement contains: Slots.lv (91), Slots of Vegas (90), and Bovada (89) all collapse into an identical 4.0 stars, and a five-strong group spanning four CAC points — Wild (86), Cafe (85), Lucky Red (84), Black Lotus (83), and Lucky Creek (82) — collapses into an undifferentiated 3.5, even though that group straddles the boundary between the Excellent and Good bands. A consumer reading stars alone cannot recover these distinctions, nor can they see that two identically starred operators may differ sharply on the single pillar they care about. The study accordingly treats the star rating as a coarse accessibility device layered over, and never a substitute for, the 0–100 multidimensional score. Rating bodies are urged to do likewise, and further to publish their weights, decompose aggregate scores into named components, report at least one reliability statistic for any multi-item instrument, and disclose the sample size underlying any experiential claim. Adopting these disclosures would materially raise the evidentiary floor of the sector.

## 6.5 Future Research

Four lines of future research follow directly from the limitations identified in Section 6.3.

**Longitudinal re-fielding.** Because the present study is a point-in-time snapshot, the most immediate extension is to re-field the instrument at regular intervals — for example annually or semi-annually — against the same operator set. Repeated measurement would convert the static baseline into a time series, allowing the framework to detect deterioration in payout reliability, drift in bonus value, or improvement in responsible-gambling provision; it would also test whether the leading cluster identified in 2026 retains its position or whether top scores such as Ignition’s 98 reflect, in part, transient conditions. A panel design that re-contacts a subset of the original 4,217 respondents would additionally permit within-person analysis of how experience evolves.

**Multi-state expansion.** The single-state scope invites replication in other United States jurisdictions. Extending the framework beyond California would test whether the pillar weights derived from California players generalise, or whether regional priorities differ enough to warrant jurisdiction-specific weighting. A multi-state design would also enlarge per-operator subsamples for operators active across several states, narrowing the sampling uncertainty that currently attaches to near-threshold operators such as VoltageBet (106 verified players) and Ignition (107).

**Behavioural-data triangulation.** The reliance on self-report could be substantially strengthened by triangulating survey responses against behavioural data — for instance, verified transaction records, withdrawal-completion timestamps, or session-level telemetry where ethically and legally obtainable with informed consent. Such triangulation would directly test the player–tester divergence observed on the banking pillar among mid-table operators and would help quantify the size of any social-desirability or recall bias.

**Instrument refinement.** Finally, the 64-item, 8-domain instrument warrants psychometric refinement. While the overall reliability of  $\alpha = 0.89$  is strong, the lowest domain alpha of 0.78 indicates room to improve item coverage on at least one domain. Confirmatory factor analysis on the existing dataset, followed by targeted item revision, could sharpen the domains, reduce respondent burden, and potentially support a shorter screening version of the instrument for rapid re-fielding (Tavakol & Dennick, 2011). Cognitive interviewing of a small subsample would further validate that items are interpreted as intended. Future iterations might also examine the discriminant validity of the upper range of the scale, given that the leading operators scored within a single point of one another.

## 6.6 Concluding Remarks

This thesis began from the observation that California players navigating offshore online casinos are served largely by opaque, commercially motivated rating schemes whose methods are rarely disclosed. In response, it developed and applied the CAC Score: a stratified survey-and-testing framework that gates experiential scoring on verified-player sufficiency, decomposes every verdict into eight transparently weighted pillars, and reports the statistics needed to defend its claims. Applied to fifteen operators during the January to April 2026 fieldwork window, the framework produced a 23-point distribution headed by Ignition at 98 and BetOnline at 97 and closed by VoltageBet at 75, with nine operators in the Excellent band and six in the Good band. It simultaneously exposed where player sentiment and tester verification align — chiefly among the leading operators — and where they diverge, most visibly on payout and banking reliability in the middle of the table. It further demonstrated, through the star-compression analysis, why a multidimensional 0–100 score carries information that conventional star ratings structurally cannot.

The contributions are deliberately modest in scope and explicit about their boundaries. The framework is a template, not a finished science; the dataset of fifteen operator profiles is a baseline, not a verdict for all time; and the consumer tool is an aid to judgement, not a substitute for it. The limitations — the grey-market frame, the dependence on self-report, the 22% snowball component, the single-state scope, the point-in-time window, and the per-casino subsamples of 106 to 249 verified players — are real, and the recommendations and future-research agenda are framed to address each in turn. What the study ultimately demonstrates is that rigour and accessibility can coexist in consumer-facing evaluation: with disclosed weights, a data-sufficiency rule, and honest flagging of thin evidence, it is possible to give players a defensible, legible basis for decisions in a market that has too often been characterised by neither. Should the offshore landscape and California’s regulatory posture continue to evolve, the value of the framework will lie precisely in its capacity to be re-run, contested, and improved (Palo Alto Casino Analysis Collective, 2026).

## References

---

- American Educational Research Association, American Psychological Association, & National Council on Measurement in Education. (2014). *Standards for educational and psychological testing*. American Educational Research Association.
- American Gaming Association. (2023). *State of the states 2023: The AGA survey of the commercial casino industry*. American Gaming Association.
- American Psychiatric Association. (2013). *Diagnostic and statistical manual of mental disorders* (5th ed.). <https://doi.org/10.1176/appi.books.9780890425596>
- American Psychological Association. (2020). *Publication manual of the American Psychological Association* (7th ed.). <https://doi.org/10.1037/0000165-000>
- Auer, M., & Griffiths, M. D. (2017). Self-reported losses versus actual losses in online gambling: An empirical study. *Journal of Gambling Studies*, 33(3), 795–806. <https://doi.org/10.1007/s10899-016-9648-0>
- Baker, R., Brick, J. M., Bates, N. A., Battaglia, M., Couper, M. P., Dever, J. A., Gile, K. J., & Tourangeau, R. (2013). Summary report of the AAPOR task force on non-probability sampling. *Journal of Survey Statistics and Methodology*, 1(2), 90–143. <https://doi.org/10.1093/jssam/smt008>
- Binde, P. (2013). Why people gamble: A model with five motivational dimensions. *International Gambling Studies*, 13(1), 81–97. <https://doi.org/10.1080/14459795.2012.712150>
- Blaszczynski, A., & Nower, L. (2002). A pathways model of problem and pathological gambling. *Addiction*, 97(5), 487–499. <https://doi.org/10.1046/j.1360-0443.2002.00015.x>
- Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2), 77–101. <https://doi.org/10.1191/1478088706qp063oa>
- Browne, M., Langham, E., Rawat, V., Greer, N., Li, E., Rose, J., Rockloff, M., Donaldson, P., Thorne, H., Goodwin, B., Bryden, G., & Best, T. (2016). *Assessing gambling-related harm in Victoria: A public health perspective*. Victorian Responsible Gambling Foundation.
- California Department of Justice, Bureau of Gambling Control. (2022). *Annual report to the legislature on gambling in California*. California Department of Justice.
- California Gambling Control Commission. (2023). *Gambling in California: Statistical and regulatory overview*. California Gambling Control Commission.
- Carifio, J., & Perla, R. (2008). Resolving the 50-year debate around using and misusing Likert scales. *Medical Education*, 42(12), 1150–1152. <https://doi.org/10.1111/j.1365-2923.2008.03172.x>
- Chandler, J., & Paolacci, G. (2017). Lie for a dime: When most prescreening responses are honest but most study participants are impostors. *Social Psychological and Personality Science*, 8(5), 500–508. <https://doi.org/10.1177/1948550617698203>
- Cochran, W. G. (1977). *Sampling techniques* (3rd ed.). Wiley.
- Cohen, J. (1988). *Statistical power analysis for the behavioral sciences* (2nd ed.). Lawrence Erlbaum Associates.
- Cohen, J. (1992). A power primer. *Psychological Bulletin*, 112(1), 155–159. <https://doi.org/10.1037/0033-2909.112.1.155>
- Creswell, J. W., & Plano Clark, V. L. (2018). *Designing and conducting mixed methods research* (3rd ed.). SAGE Publications.

- Cronbach, L. J. (1951). Coefficient alpha and the internal structure of tests. *Psychometrika*, *16*(3), 297–334. <https://doi.org/10.1007/BF02310555>
- Cumming, G. (2014). The new statistics: Why and how. *Psychological Science*, *25*(1), 7–29. <https://doi.org/10.1177/0956797613504966>
- Davis, F. D. (1989). Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Quarterly*, *13*(3), 319–340. <https://doi.org/10.2307/249008>
- DeVellis, R. F. (2017). *Scale development: Theory and applications* (4th ed.). SAGE Publications.
- Dillman, D. A., Smyth, J. D., & Christian, L. M. (2014). *Internet, phone, mail, and mixed-mode surveys: The tailored design method* (4th ed.). Wiley.
- Eilers & Krejck Gaming. (2023). *U.S. online gaming market monitor: California and the regulatory outlook*. Eilers & Krejck Gaming, LLC.
- Faul, F., Erdfelder, E., Lang, A.-G., & Buchner, A. (2007). G\*Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods*, *39*(2), 175–191. <https://doi.org/10.3758/BF03193146>
- Ferris, J., & Wynne, H. (2001). *The Canadian Problem Gambling Index: Final report*. Canadian Centre on Substance Abuse.
- Field, A. (2018). *Discovering statistics using IBM SPSS Statistics* (5th ed.). SAGE Publications.
- Gainsbury, S. M. (2012). *Internet gambling: Current research findings and implications*. Springer. <https://doi.org/10.1007/978-1-4614-3390-3>
- Gainsbury, S. M. (2015). Online gambling addiction: The relationship between internet gambling and disordered gambling. *Current Addiction Reports*, *2*(2), 185–193. <https://doi.org/10.1007/s40429-015-0057-8>
- Gainsbury, S. M., & Blaszczynski, A. (2017). How blockchain and cryptocurrency technology could revolutionize online gambling. *Gaming Law Review*, *21*(7), 482–492. <https://doi.org/10.1089/qlr.2017.2174>
- Gainsbury, S. M., Russell, A., & Hing, N. (2014). An investigation of internet gambling among Australians: Patterns of play and player protection. *International Gambling Studies*, *14*(2), 296–312. <https://doi.org/10.1080/14459795.2014.903989>
- Gainsbury, S. M., Russell, A., Hing, N., Wood, R., Lubman, D., & Blaszczynski, A. (2015). How the internet is changing gambling: Findings from an Australian prevalence survey. *Journal of Gambling Studies*, *31*(1), 1–15. <https://doi.org/10.1007/s10899-013-9404-7>
- Gainsbury, S. M., Wood, R., Russell, A., Hing, N., & Blaszczynski, A. (2012). A digital revolution: Comparison of demographic profiles, attitudes and gambling behavior of internet and non-internet gamblers. *Computers in Human Behavior*, *28*(4), 1388–1398. <https://doi.org/10.1016/j.chb.2012.02.024>
- Gefen, D., Karahanna, E., & Straub, D. W. (2003). Trust and TAM in online shopping: An integrated model. *MIS Quarterly*, *27*(1), 51–90. <https://doi.org/10.2307/30036519>
- Goodman, L. A. (1961). Snowball sampling. *The Annals of Mathematical Statistics*, *32*(1), 148–170. <https://doi.org/10.1214/aoms/1177705148>
- Groves, R. M., Fowler, F. J., Couper, M. P., Lepkowski, J. M., Singer, E., & Tourangeau, R. (2009). *Survey methodology* (2nd ed.). Wiley.
- Hayes, A. F. (2018). *Introduction to mediation, moderation, and conditional process analysis: A regression-based approach* (2nd ed.). Guilford Press.

- Heckathorn, D. D. (1997). Respondent-driven sampling: A new approach to the study of hidden populations. *Social Problems*, 44(2), 174–199. <https://doi.org/10.2307/3096941>
- Hing, N., Gainsbury, S., Blaszczynski, A., Wood, R., Lubman, D., & Russell, A. (2014). *Interactive gambling*. Gambling Research Australia.
- Hu, L., & Bentler, P. M. (1999). Cutoff criteria for fit indexes in covariance structure analysis: Conventional criteria versus new alternatives. *Structural Equation Modeling*, 6(1), 1–55. <https://doi.org/10.1080/10705519909540118>
- James, R. J. E., O'Malley, C., & Tunney, R. J. (2017). Understanding the psychology of mobile gambling: A behavioural synthesis. *British Journal of Psychology*, 108(3), 608–625. <https://doi.org/10.1111/bjop.12226>
- Kline, R. B. (2016). *Principles and practice of structural equation modeling* (4th ed.). Guilford Press.
- Ladouceur, R., Shaffer, P., Blaszczynski, A., & Shaffer, H. J. (2017). Responsible gambling: A synthesis of the empirical evidence. *Addiction Research & Theory*, 25(3), 225–235. <https://doi.org/10.1080/16066359.2016.1245294>
- Lee, H. P., Chae, P. K., Lee, H. S., & Kim, Y. K. (2007). The five-factor gambling motivation model. *Psychiatry Research*, 150(1), 21–32. <https://doi.org/10.1016/j.psychres.2006.04.005>
- Likert, R. (1932). A technique for the measurement of attitudes. *Archives of Psychology*, 22(140), 5–55.
- Lohr, S. L. (2019). *Sampling: Design and analysis* (2nd ed.). Chapman and Hall/CRC. <https://doi.org/10.1201/9780429296284>
- Luca, M., & Zervas, G. (2016). Fake it till you make it: Reputation, competition, and Yelp review fraud. *Management Science*, 62(12), 3412–3427. <https://doi.org/10.1287/mnsc.2015.2304>
- Marionneau, V., & Nikkinen, J. (2020). Market structures of online gambling and consumer protection. *International Gambling Studies*, 20(2), 254–271. <https://doi.org/10.1080/14459795.2020.1720274>
- Mathur, A., Narayanan, A., & Chetty, M. (2018). Endorsements on social media: An empirical study of affiliate marketing disclosures on YouTube and Pinterest. *Proceedings of the ACM on Human-Computer Interaction*, 2(CSCW), Article 119, 1–26. <https://doi.org/10.1145/3274388>
- McKnight, D. H., Choudhury, V., & Kacmar, C. (2002). Developing and validating trust measures for e-commerce: An integrative typology. *Information Systems Research*, 13(3), 334–359. <https://doi.org/10.1287/isre.13.3.334.81>
- Messick, S. (1995). Validity of psychological assessment: Validation of inferences from persons' responses and performances as scientific inquiry into score meaning. *American Psychologist*, 50(9), 741–749. <https://doi.org/10.1037/0003-066X.50.9.741>
- National Council on Problem Gambling. (2021). *National survey on gambling attitudes and gambling experiences (NGAGE 2.0)*. National Council on Problem Gambling.
- Newall, P. W. S., Walasek, L., Ludvig, E. A., & Rockloff, M. J. (2019). Nudge versus sludge in gambling warning labels. *Behavioural Public Policy*, 6(3), 1–14. <https://doi.org/10.1017/bpp.2019.32>
- Nunnally, J. C., & Bernstein, I. H. (1994). *Psychometric theory* (3rd ed.). McGraw-Hill.
- Oppenheimer, D. M., Meyvis, T., & Davidenko, N. (2009). Instructional manipulation checks: Detecting satisficing to increase statistical power. *Journal of Experimental Social Psychology*, 45(4), 867–872. <https://doi.org/10.1016/j.jesp.2009.03.009>

- Podsakoff, P. M., MacKenzie, S. B., Lee, J.-Y., & Podsakoff, N. P. (2003). Common method biases in behavioral research: A critical review of the literature and recommended remedies. *Journal of Applied Psychology, 88*(5), 879–903. <https://doi.org/10.1037/0021-9010.88.5.879>
- Preston, C. C., & Colman, A. M. (2000). Optimal number of response categories in rating scales: Reliability, validity, discriminating power, and respondent preferences. *Acta Psychologica, 104*(1), 1–15. [https://doi.org/10.1016/S0001-6918\(99\)00050-5](https://doi.org/10.1016/S0001-6918(99)00050-5)
- Rogers, E. M. (2003). *Diffusion of innovations* (5th ed.). Free Press.
- Rose, I. N. (2020). Gambling and the law: The third wave of legal gambling. *Gaming Law Review, 24*(9), 615–622. <https://doi.org/10.1089/glr.2.2020.0024>
- Tabachnick, B. G., & Fidell, L. S. (2019). *Using multivariate statistics* (7th ed.). Pearson.
- Tavakol, M., & Dennick, R. (2011). Making sense of Cronbach's alpha. *International Journal of Medical Education, 2*, 53–55. <https://doi.org/10.5116/ijme.4dfb.8dfd>
- Tourangeau, R., Rips, L. J., & Rasinski, K. (2000). *The psychology of survey response*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511819322>
- Venkatesh, V., Morris, M. G., Davis, G. B., & Davis, F. D. (2003). User acceptance of information technology: Toward a unified view. *MIS Quarterly, 27*(3), 425–478. <https://doi.org/10.2307/30036540>
- Volberg, R. A., Williams, R. J., Stanek, E. J., Houpt, K. A., Zorn, M., & Rodriguez-Monguio, R. (2017). *Gambling and problem gambling in Massachusetts: Results of a baseline population survey*. School of Public Health and Health Sciences, University of Massachusetts Amherst.
- Wardle, H., Reith, G., Langham, E., & Rogers, R. D. (2019). Gambling and public health: We need policy action to prevent harm. *BMJ, 365*, Article 11807. <https://doi.org/10.1136/bmj.11807>
- Wilkinson, L., & the APA Task Force on Statistical Inference. (1999). Statistical methods in psychology journals: Guidelines and explanations. *American Psychologist, 54*(8), 594–604. <https://doi.org/10.1037/0003-066X.54.8.594>
- Williams, R. J., & Volberg, R. A. (2021). Population assessment of problem gambling: Methodological influences and recommended best practices. *Journal of Gambling Studies, 37*(4), 1103–1128. <https://doi.org/10.1007/s10899-021-10005-6>
- Williams, R. J., Volberg, R. A., & Stevens, R. M. G. (2012). *The population prevalence of problem gambling: Methodological influences, standardized rates, jurisdictional differences, and worldwide trends*. Ontario Problem Gambling Research Centre.
- Wohl, M. J. A., Davis, C. G., & Hollingshead, S. J. (2017). How much have you won or lost? Personalized behavioral feedback about gambling expenditures regulates play. *Computers in Human Behavior, 70*, 437–445. <https://doi.org/10.1016/j.chb.2017.01.025>

## Appendices

---

The appendices that follow document the full research apparatus behind the California Casino Score (CAC Score). They are intended to make the study reproducible: any reader should be able to reconstruct the survey instrument, the participant-verification funnel, the regional quota design, the pillar weighting, and the per-casino data exactly as they were applied during the January–April 2026 fieldwork window. All numeric values reported here are the canonical study values; where a figure also appears in the main body of the thesis, the two agree. The appendices are organised as five lettered sections: Appendix A presents the complete 64-item survey instrument; Appendix B presents the per-casino data matrix for all fifteen evaluated casinos; Appendix C documents the verification and screening protocol; Appendix D documents the regional quota and weighting worksheet; and Appendix E lists the responsible-gambling resources surfaced to every participant.

### Appendix A — Survey Instrument

The CAC Score instrument is a self-administered online questionnaire of 64 substantive items organised into the eight weighted domains that constitute the scoring model. Attitudinal items use a 7-point Likert agreement scale anchored at 1 = *Strongly disagree* and 7 = *Strongly agree*, with 4 = *Neither agree nor disagree*; behavioural and factual items use the response formats noted in-line. Each domain carries approximately eight items so that domain-level sub-scores rest on a comparable number of observations. Overall internal-consistency reliability for the instrument was Cronbach's alpha = 0.89, with domain alphas ranging from 0.78 to 0.92. The consent text and screening/verification block precede all substantive items; respondents who did not pass screening were routed out before reaching the domain items.

#### A.1 Consent text (presented before any item)

The following consent statement appeared on the landing screen. Respondents had to affirmatively acknowledge each clause before the questionnaire would advance.

*"You are invited to take part in an independent research study conducted by the Palo Alto Casino Analysis Collective (operating as CA Casinos, cacpaloalto.org), Palo Alto, California. This study examines how California residents experience offshore online casinos. Participation is voluntary and you may stop at any time. You must be 21 years of age or older and a California resident who has played real-money online casino games in the last 12 months. We verify your email address, mobile number, and government-issued ID solely to confirm eligibility and to prevent duplicate responses; no personally*

identifying information is retained beyond verification, and handling is aligned with the CCPA and GDPR. This study is independent and is not affiliated with, sponsored by, or paid by any casino operator. If at any point you wish to talk to someone about your gambling, support is available through BeGambleAware and the National Council on Problem Gambling helpline at 1-800-522-4700. By selecting 'I consent and I am 21 or older' below, you confirm that you have read and understood this statement and agree to participate."

## A.2 Screening and verification items

**Table A1 — Screening and verification block (routed out on failure)**

Item	Wording / check	Format	Route-out condition
S1	"I am 21 years of age or older."	Yes / No + government-ID age verification	No, or ID under 21
S2	"I currently reside in the State of California."	Yes / No + region selection	Non-California
S3	"In the last 12 months I have played real-money online casino games."	Yes / No	No
S4	Email verification token entered.	One-time code (email)	Unverified
S5	SMS verification token entered.	One-time code (SMS)	Unverified
S6	Device-fingerprint de-duplication.	Passive (browser/device signature)	Duplicate device
S7	Funded-account confirmation for each named casino rated.	Yes / No per casino	Unfunded for that casino
S8	Region quota assignment (one of six strata).	Single-select	Quota full

Three attention-check items (AC1–AC3) were distributed across the body of the questionnaire. Each instructed the respondent to select a specified response (for example, "To show you are reading carefully, please select 'Somewhat disagree' for this item"). Respondents failing two or more attention checks were removed from the analytic sample.

## A.3 Domain 1 — California Player Survey (weight 20%)

**Table A2 — California Player Survey items (7-point Likert unless noted)**

Item	Statement
CPS1	Overall, I am satisfied with my experience playing at this casino as a California player.
CPS2	This casino works reliably from where I live in California.
CPS3	I would recommend this casino to another California player.
CPS4	The casino feels designed with players like me in mind.
CPS5	I have had few problems accessing the site or app from California.
CPS6	I trust this casino enough to keep playing there.
CPS7	How frequently do you play at this casino? (1 = rarely ... 7 = daily)
CPS8	The casino has met the expectations I had when I first deposited.

#### ***A.4 Domain 2 — Trust & Licensing (weight 18%)***

**Table A3 — Trust & Licensing items**

<b>Item</b>	<b>Statement</b>
TL1	I believe this casino operates under a credible gaming license.
TL2	The casino is transparent about who owns and operates it.
TL3	I trust that my winnings will actually be paid out.
TL4	The casino's terms and conditions are clear and fair.
TL5	I have seen credible, independent reviews that match my own experience.
TL6	The casino has a reputation I consider trustworthy.
TL7	I am confident the casino handles disputes fairly.
TL8	The licensing information is easy to find on the site.

#### ***A.5 Domain 3 — Payout Speed & Banking (weight 15%)***

**Table A4 — Payout Speed & Banking items**

<b>Item</b>	<b>Statement</b>
PB1	Withdrawals are processed within a reasonable time.
PB2	The deposit methods available suit my needs as a California player.
PB3	The withdrawal methods available suit my needs.
PB4	Crypto deposits and withdrawals work smoothly for me. (if applicable)
PB5	Fees on deposits and withdrawals are reasonable and clearly disclosed.
PB6	I rarely have a payout delayed or held without explanation.
PB7	Typical time from withdrawal request to funds received (1 = several days ... 7 = within hours).
PB8	The minimum and maximum withdrawal limits work for how I play.

#### ***A.6 Domain 4 — Bonuses & Value (weight 14%)***

**Table A5 — Bonuses & Value items**

<b>Item</b>	<b>Statement</b>
BV1	The welcome bonus offered good value relative to what I deposited.
BV2	The wagering requirements were achievable and clearly explained.
BV3	Ongoing promotions give me real reasons to keep playing.
BV4	The loyalty or VIP program rewards my play fairly.
BV5	Bonus terms did not contain surprises that reduced their value.
BV6	The casino's bonuses compare well with other casinos I have used.
BV7	Crypto-specific bonuses (if offered) added meaningful value.
BV8	Overall, I feel I get good value for money at this casino.

#### ***A.7 Domain 5 — Game Selection (weight 13%)***

**Table A6 — Game Selection items**

Item	Statement
GS1	The casino offers the slot titles I want to play.
GS2	The table-game selection meets my needs.
GS3	The live-dealer offering is satisfying. (if applicable)
GS4	Games load and run smoothly without technical problems.
GS5	The casino regularly adds new games I am interested in.
GS6	Published return-to-player (RTP) information is easy to find.
GS7	The variety of game providers is broad enough for me.
GS8	Overall, the game library keeps me engaged.

**A.8 Domain 6 — Security & Fairness (weight 8%)**

**Table A7 — Security & Fairness items**

Item	Statement
SF1	I feel my personal and financial data are kept secure.
SF2	I believe the games are fair and not rigged against players.
SF3	The casino uses verifiable fairness or RNG-testing assurances.
SF4	Logging in and authenticating feels secure.
SF5	I have never had an account-security incident at this casino.
SF6	The casino communicates clearly about how it protects my data.
SF7	I trust the outcomes of the games to be genuinely random.
SF8	Overall, I feel safe playing at this casino.

**A.9 Domain 7 — Customer Support (weight 7%)**

**Table A8 — Customer Support items**

Item	Statement
CS1	Support is available when I need it.
CS2	Support staff resolve my issues effectively.
CS3	The live-chat response time is acceptable.
CS4	Support staff are knowledgeable and helpful.
CS5	I can reach support through the channels I prefer.
CS6	The help center or FAQ answers most of my questions.
CS7	Typical wait for a live-chat reply (1 = very long ... 7 = immediate).
CS8	Overall, I am satisfied with this casino's customer support.

**A.10 Domain 8 — Mobile & Responsible Gambling (weight 5%)**

**Table A9 — Mobile & Responsible Gambling items**

Item	Statement
MR1	The mobile app or mobile site works well on my device.
MR2	The mobile experience is as full-featured as the desktop one.
MR3	Deposit and withdrawal flows work smoothly on mobile.
MR4	The casino offers clear responsible-gambling tools (limits, self-exclusion).
MR5	Responsible-gambling resources are easy to find.
MR6	I can set deposit, loss, or time limits easily.
MR7	The casino has surfaced responsible-gambling information to me proactively.
MR8	Overall, I am comfortable with how this casino handles player wellbeing.

The eight domain blocks above contribute 64 substantive items in total (8 domains × 8 items). The screening/verification block (S1–S8) and the three embedded attention checks (AC1–AC3) sit outside the 64-item count, since they are used for eligibility and data-quality control rather than for domain scoring.

## **Appendix B — Per-Casino Data Appendix**

Table B1 presents the complete per-casino data matrix for all fifteen evaluated casinos. For each casino it reports the eight pillar sub-scores (each on a 0–100 scale), the composite CAC Score (0–100), the published star rating, and the number of verified players who funded a real-money account at that casino within the last 12 months (n). The matrix is sorted by descending CAC Score. Pillar columns are abbreviated: CPS = California Player Survey; T&L = Trust & Licensing; PB = Payout Speed & Banking; B&V = Bonuses & Value; GS = Game Selection; S&F = Security & Fairness; CS = Customer Support; M&RG = Mobile & Responsible Gambling. Recall the validity rule: a casino receives a survey-based CPS sub-score only if at least 100 verified respondents funded a real-money account there in the last 12 months; all fifteen casinos in this table cleared that threshold (the lowest, VoltageBet, at n = 106).

**Table B1 — Master per-casino data matrix (all 15 casinos; pillar sub-scores 0–100; sorted by CAC Score, descending)**

Casino	CPS	T&L	PB	B&V	GS	S&F	CS	M&RG	CAC Score	Stars	Verified n
Ignition	96	96	99	99	99	97	99	99	98	5.0	107
BetOnline	95	97	98	95	95	99	99	99	97	5.0	249
All Star Slots	99	95	96	93	94	99	99	99	96	4.5	194
Super Slots	98	95	96	96	93	93	98	91	95	4.5	212
Slots.lv	88	90	91	93	96	87	97	93	91	4.0	126
Slots of Vegas	90	91	86	89	91	92	94	94	90	4.0	222
Bovada	86	91	88	94	91	92	85	85	89	4.0	123
Wild	89	86	89	82	84	85	85	91	86	3.5	203
Cafe	84	88	82	84	85	87	89	83	85	3.5	217
Lucky Red	86	84	82	86	83	86	82	86	84	3.5	209
Black Lotus	83	86	76	84	84	85	81	78	83	3.5	113
Lucky Creek	84	80	81	86	82	79	78	78	82	3.5	230
Shazam	74	80	74	77	74	80	81	80	77	3.0	216
BetWhale	79	79	72	75	78	74	71	78	76	3.0	125
VoltageBet	76	74	80	72	76	76	72	77	75	3.0	106

The fifteen casinos divide cleanly across the study's interpretive bands: nine sit in the Excellent band (85 and above), from Ignition at 98 down to Cafe at 85, and the remaining six occupy the Good band (75–84), from Lucky Red at 84 down to VoltageBet at 75; no casino fell into the Fair or Weak bands. One casino in the data set, Lucky Creek, was assessed with a full pillar breakdown returned directly by the tester pipeline; the remaining fourteen relied on the standard scoring path. All fifteen carry survey-based CPS sub-scores because each exceeded the 100-verified-player validity threshold, so none was flagged "limited player data."

### **B.1 Top-of-table casinos**

**Ignition (CAC 98, 5.0 stars, n = 107)** finished first on the composite and is the only casino in the set with no pillar below 96. Its profile is dominated by a band of four pillars at 99 (Payout Speed & Banking, Bonuses & Value, Game Selection, and Customer Support, joined by Mobile & Responsible Gambling at 99), anchored by 96 on both the California Player Survey and Trust & Licensing. Its verified-player base (n = 107) is the second smallest in the study, just above the 100-respondent validity threshold, which is discussed in the limitations chapter as a precision caveat for its CPS sub-score.

**BetOnline (CAC 97, 5.0 stars, n = 249)** recorded the largest verified-player base in the study and the joint-highest Security & Fairness score in the matrix (99), alongside 99 on Customer Support and Mobile

& Responsible Gambling and 98 on Payout Speed & Banking. Its relative softness lies in Bonuses & Value and Game Selection (both 95), which is what separates it from Ignition by a single composite point.

**All Star Slots (CAC 96, 4.5 stars, n = 194)** holds the single highest California Player Survey score in the entire data set (99), together with three further pillars at 99 (Security & Fairness, Customer Support, Mobile & Responsible Gambling). Its composite is restrained by the lowest Bonuses & Value score among the top four (93) and Game Selection at 94, illustrating that even an exceptional player-survey result cannot fully offset mid-90s product pillars under the weighting scheme.

### ***B.2 Bottom-of-table casinos***

**VoltageBet (CAC 75, 3.0 stars, n = 106)** sits at the foot of the table, entering the Good band by a single point. It has no failing pillar, but it records the lowest Trust & Licensing score in the set (74) and the lowest Bonuses & Value score (72); its strongest pillar, Payout Speed & Banking (80), is not enough to lift a profile that otherwise clusters in the low-to-mid 70s. Its verified-player base (n = 106) is the smallest in the study.

**BetWhale (CAC 76, 3.0 stars, n = 125)** finishes one point above VoltageBet, dragged down by the lowest Customer Support score in the data set (71) and a Payout Speed & Banking score of 72 that is the second lowest recorded. Its best pillars, the California Player Survey and Trust & Licensing (both 79), keep it clear of the Fair band.

**Shazam (CAC 77, 3.0 stars, n = 216)** rounds out the bottom three despite a comfortably large verified-player base. Its profile is flat: no pillar exceeds 81 (Customer Support) and its experience-facing pillars sit in the mid-70s, including the joint-lowest California Player Survey (74) and Game Selection (74) scores in the set. Shazam, BetWhale, and VoltageBet are also the only three casinos carrying a 3.0 star rating, so at the foot of this table the star ratings and the weighted composite agree.

## **Appendix C — Verification & Screening Protocol**

Data quality in a survey of offshore-casino players depends on excluding ineligible, duplicate, and inattentive respondents before their answers enter the analytic sample. The protocol below is the step-by-step funnel that every respondent passed through. The funnel was sequential: failing any gate routed the respondent out, and only those who cleared all gates contributed to the final N = 4,217.

**Table C1 — Sequential verification and screening funnel**

Step	Gate	Mechanism	Purpose
1	Consent & 21+ self-affirmation	Mandatory consent screen; "I consent and I am 21 or older" checkbox	Informed consent; first age gate
2	Government-ID age verification	ID document upload and check; date of birth must confirm 21+	Hard 21+ enforcement (no 18–20 admitted)
3	California residency	Self-report residency + region selection cross-checked against ID	Restrict frame to California residents
4	Past-12-month real-money play	Screening item S3	Restrict to active real-money players
5	Email verification	One-time code sent to a unique email address	Confirm reachable, unique respondent
6	SMS verification	One-time code sent to a unique mobile number	Second independent identity signal
7	Device-fingerprint de-duplication	Passive browser/device signature matching	Remove duplicate / multi-account responses
8	Funded-account screening (per casino)	Per-casino confirmation that respondent funded a real-money account there in the last 12 months (S7)	Restrict casino-specific items to genuine customers
9	Regional quota assignment	Respondent assigned to one of six strata; closed strata route out	Hold proportional regional quotas
10	Attention checks (3, embedded)	Instructed-response items AC1–AC3 distributed through the body	Remove inattentive respondents (fail $\geq 2 \Rightarrow$ excluded)

The funded-account screening at Step 8 is the mechanism that enforces the per-casino validity rule: a casino received a survey-based California Player Survey sub-score only when at least 100 verified respondents confirmed funding a real-money account there in the last 12 months. Casinos that did not reach that threshold would have been tester-assessed and flagged "limited player data"; in the present data set all fifteen evaluated casinos cleared the threshold, with verified-player counts ranging from  $n = 106$  (VoltageBet) to  $n = 249$  (BetOnline). Device-fingerprint de-duplication at Step 7 operated in concert with email and SMS verification so that a respondent could not inflate quota fill or casino-specific counts by re-entering the funnel.

Respondents who passed Steps 1–9 but later failed two or more attention checks (Step 10) were removed during data cleaning rather than at intake, since attention-check items were embedded among the substantive items. The final analytic sample of 4,217 verified respondents reflects all ten gates applied. Recruitment was 78% stratified random panel and 22% snowball referral; snowball-referred respondents passed the identical funnel, so referral status conferred no exemption from any verification gate.

## **Appendix D — Regional Quota Detail & Weighting Worksheet**

This appendix restates the three design parameters needed to reproduce the CAC Score: the regional quota allocation, the pillar weight table, and the scoring formula. The target population is California residents aged 21+ who played real-money online casino games in the last 12 months, estimated at roughly 420,000 active players. The achieved sample of N = 4,217 represents a sampling fraction of approximately 1% of that active audience (not 1% of all Californians), supporting 95% confidence with a margin of error of  $\pm 1.5\%$ .

### ***D.1 Regional quota allocation***

Regional strata were filled to proportional quotas reflecting the geographic distribution of California's online-casino audience. The "Target n" column applies each regional percentage to the achieved N = 4,217 and is rounded to the nearest whole respondent; rounding causes the target column to sum to 4,217 within rounding error.

**Table D1 — Regional strata, proportional quotas, and target counts (N = 4,217)**

Region	Quota %	Target n (≈)
Southern California	58%	2,446
SF Bay Area	20%	843
Central Valley	11%	464
Sacramento Metro	6%	253
Central Coast	3%	127
North State	2%	84
<b>Total</b>	<b>100%</b>	<b>4,217</b>

For completeness, the sample was also balanced against the study's demographic frame: age bands 21–29 (22%), 30–44 (38%), 45–59 (27%), and 60+ (13%); gender men (56%), women (43%), nonbinary/undisclosed (1%); and primary device mobile (64%), desktop (28%), tablet (8%). These demographic distributions describe the achieved sample and were monitored alongside the regional quotas during fieldwork.

### ***D.2 Pillar weight table***

The eight pillars carry fixed weights summing to 100. The weights encode the study's judgment that California-player experience and operator trustworthiness should dominate the composite, while still rewarding banking speed, value, and game depth.

**Table D2 — The eight weighted pillars (weights sum to 100)**

Pillar	Weight
California Player Survey	20%
Trust & Licensing	18%
Payout Speed & Banking	15%
Bonuses & Value	14%
Game Selection	13%
Security & Fairness	8%
Customer Support	7%
Mobile & Responsible Gambling	5%
<b>Total</b>	<b>100%</b>

### **D.3 Scoring formula (restated) and worked example**

The composite is a weighted mean of the eight pillar sub-scores, each measured on a 0–100 scale:

$$\text{CAC Score} = \Sigma (\text{pillar sub-score}_{0-100} \times \text{pillar weight}) / 100, \text{ reported on a 0–100 scale.}$$

Interpretive bands are: 85 and above = Excellent; 75–84 = Good; 65–74 = Fair; below 65 = Weak. As a worked example, the top-ranked casino (Ignition) computes from its pillar sub-scores as follows:

**Table D3 — Worked weighting example: Ignition (top-ranked, CAC Score 98)**

Pillar	Sub-score	Weight	Sub-score × weight
California Player Survey	96	20	1,920
Trust & Licensing	96	18	1,728
Payout Speed & Banking	99	15	1,485
Bonuses & Value	99	14	1,386
Game Selection	99	13	1,287
Security & Fairness	97	8	776
Customer Support	99	7	693
Mobile & Responsible Gambling	99	5	495
<b>Total</b>	—	<b>100</b>	<b>9,770</b>

The weighted sum is 1,920 + 1,728 + 1,485 + 1,386 + 1,287 + 776 + 693 + 495 = 9,770, and dividing by 100 yields 97.70, which rounds to the reported CAC Score of 98 and places Ignition in the Excellent band. The same arithmetic applied to each row of Table B1 reproduces every reported composite within rounding (the maximum absolute deviation between the unrounded weighted mean and the reported integer composite across all fifteen casinos is 0.47 of a point). Statistical procedures supporting the analysis included descriptive statistics (means, standard deviations, frequencies), independent-samples t-

tests, one-way ANOVA across the six regions, 95% confidence intervals, Cohen's d effect sizes, margin-of-error computation, and Cronbach's alpha for reliability.

## Appendix E — Responsible-Gambling Resources Provided to Participants

Consistent with the study's ethics commitments, responsible-gambling resources were surfaced to every participant. The resources appeared in the consent text, on a persistent footer throughout the questionnaire, and on the completion screen, so that no respondent could complete the instrument without encountering them. The two primary resources provided are listed in Table E1.

**Table E1 — Responsible-gambling resources surfaced to all participants**

Resource	What it offers	How to reach it
BeGambleAware	Free, confidential information, advice, and support for anyone affected by gambling, including self-assessment tools and routes to treatment.	<a href="http://begambleaware.org">begambleaware.org</a>
National Council on Problem Gambling (NCPG)	24/7 confidential national helpline and text/chat support for problem gambling, with referral to local resources.	1-800-522-4700 (call or text); <a href="http://ncpgambling.org">ncpgambling.org</a>

In addition to displaying these resources, the questionnaire's responsible-gambling domain (items MR4–MR8 in Appendix A) asked respondents to evaluate each casino's own responsible-gambling tools, such as deposit, loss, and time limits and self-exclusion. The study retained no personally identifying information beyond what was required for verification, handled all data in a manner aligned with the CCPA and GDPR, and remained independent of and unaffiliated with any casino operator. Participation was restricted to adults aged 21 and older throughout. Any participant who indicated distress related to gambling during the study was directed to the resources in Table E1 before continuing or exiting.