How ThumbAd Answers the 20 Questions

We published a due-diligence guide telling iGaming operators exactly what to demand from a programmatic partner. It would be strange to duck our own exam. Here is every question, answered in writing — with the artifacts we bring to prove it in person.

```
> load igos/questions.20                [ OK ]
> mode: written answers, on the record[ OK ]
> evasion.dll                    [ NOT FOUND ]
```

Q01-06

Data & Audiences

Ownership, behaviour-built segments, consent, separation.

Q07-13

Measurement & Proof

Path data, reconciliation, incrementality, honest KPIs.

Q14-20

Money, Ops & Fit

Fees, incentives, people, compliance, and one lost account.

Why we're answering our own exam

The companion guide — *The iGaming Operator's Guide to Programmatic* — argues that the gap between confident and evasive answers is the whole due-diligence signal. Publishing our answers is therefore not bravado; it's the logical consequence. If the questions are fair, they're fair for us.

Three ways to use this document. Read it before meeting us, and the first meeting starts three meetings deep. Bring it to the meeting and score us live against what's printed — we'll bring the artifacts each answer references. Or use it against your current partner: every answer here is a benchmark for what the same question should produce elsewhere.

// WHAT'S PRINTED HERE	// WHAT WE BRING IN PERSON
Q07 — yes to raw path data	An anonymised path-to-conversion export
Q10 — yes, we cut good-looking channels	The specific story: client, channel, data, aftermath
Q14 — every fee, line by line	A sample fee sheet you can keep
Q18 — decisions, not metrics dumps	Real weekly, monthly and quarterly examples
Q20 — a lost account, honestly	The account, the reason, the process change

One reading note: where an answer depends on a specific contract, platform structure or market, we say so in the answer. Certainty we don't have is not something we print — that discipline is also part of the answer.

DATA & AUDIENCES

Q01

Who owns the pixel and the audiences you build?

// THE SHORT ANSWER

You do. In your seat, documented, portable.

// IN PRACTICE

Client ownership of pixels, audiences and campaign learning is our default contract position, not a concession you have to negotiate. Segments are built in your advertiser structure, named to a documented taxonomy, and written up so they make sense without us in the room.

Where a platform's seat structure or privacy rules limit what can physically transfer, we put that in writing before signing — you'll know exactly what leaves with you and what is platform-bound. What we will never do is treat your audience intelligence as leverage against the day you want to leave.

HOLD US TO IT Ask for the audience taxonomy document in month one. If it doesn't exist, this answer was theatre.

DATA & AUDIENCES

Show me an audience you built from behaviour, not demographics.

Q02

// THE SHORT ANSWER

Here’s one: started registration, never finished.

// IN PRACTICE

That single behaviour outperforms any demographic label, because it's a person telling you exactly where they stopped. Our iGaming segments are built from events like it:

> Registration started, not completed

> Registered, never deposited

> Lapsed depositors flagged for reactivation from CRM

> High-intent visitors by funnel depth, frequency and recency

> Product- and market-page behaviour, not “interest in sports”

> Exposed on CTV or DOOH, reachable again on mobile and display

“Males 25–44, sports interest” is a census category, not an audience. We build from what people did.

HOLD US TO IT Pick any live segment in your account and ask us to show the behaviours that define it.

DATA & AUDIENCES

How do you prove a new audience beats what it replaces?

Q03

// THE SHORT ANSWER

A named test, designed in writing before launch.

// IN PRACTICE

Every challenger audience gets a test card before it spends: the incumbent it must beat, the success metric, the budget split, the runtime, the minimum conversion volume, the attribution window, the exclusion logic, and the decision rule — scale, iterate or kill.

The metric matches your economics, not our convenience: FTD or deposit value where backend data is connected, and where it isn't yet, results are labelled platform-attributed — never dressed up as confirmed customers.

HOLD US TO IT Ask to see the test card before the test, and the verdict after. Both should exist as documents.

DATA & AUDIENCES

How do you handle consent and cookieless markets?

Q04

// THE SHORT ANSWER

Segment by segment, market by market — with consent status documented.

// IN PRACTICE

“The DSP handles it” is not an answer, so we don't give it. Every segment carries its legal basis and consent status in its documentation, and addressability is planned per market.

> First-party and consented CRM data where the basis is clean

> Contextual and inventory-based targeting where identity is thin

> UID2 / EUID and publisher audiences where supported and lawful

> Measurement expectations adjusted honestly — not pretended around

In restricted markets we change the strategy and say so, rather than running the same tactics and hoping the reporting hides it.

HOLD US TO IT Ask for the consent status of any segment we run. The answer should take minutes, not a legal review.

DATA & AUDIENCES

What happens to my data if we part ways?

Q05

// THE SHORT ANSWER

It leaves with you — and the map of what transfers is agreed before we start.

// IN PRACTICE

Setup, naming conventions, audience definitions, reporting logic and campaign learning are documented from day one, which is what makes them portable. Before signing, we walk through what is contractually and technically transferable versus what any platform keeps regardless of partner.

No hidden lock-in dressed up as partnership. A desk that needs captivity to keep clients is telling you what its work is worth.

HOLD US TO IT Ask for the offboarding clause before you need it. Ours is written to be read.

DATA & AUDIENCES

How is my data separated from your other clients?

Q06

// THE SHORT ANSWER

Client-scoped structures. No cross-account mixing, ever.

// IN PRACTICE

Your data lives in your campaigns, your advertiser structures and your projects — separated at the account level of every platform we touch, from The Trade Desk to measurement and reporting tools.

In a regulated category like iGaming this is more than hygiene: a segment that has quietly absorbed another brand's data is a compliance incident waiting to be discovered. Client-specific data stays client-specific, without exceptions for convenience.

HOLD US TO IT Ask us to diagram your account structure and where its boundaries sit. It's a one-page answer.

MEASUREMENT & PROOF

Can I see raw path-to-conversion data for my campaigns?

Q07

// THE SHORT ANSWER

Yes — and we bring an example export to the first meeting.

WE BRING: ANONYMISED PTC EXPORT

// IN PRACTICE

Path-to-conversion analysis is core to how we work, not a special request: full journey sequences with every touch, channel, creative and device, first and last exposure, touch counts, time to conversion, and view-through versus click-through roles.

Where privacy rules or platform policy limit granularity in a market, the export says so in its notes — the limitation is disclosed, not smoothed over. You will never be asked to accept screenshots as data.

HOLD US TO IT Ask for the sample export in the sales process and time the response.

MEASUREMENT & PROOF

Q08

Show me last-touch and multi-touch side by side for one campaign.

// THE SHORT ANSWER

Side by side in every review — with the disagreement explained, not hidden.

// IN PRACTICE

The two models answer different questions, so of course they disagree: last-touch flatters retargeting and lower-funnel display; multi-touch reveals what CTV, video and prospecting actually contributed upstream. The interesting part is the gap between them — that's where budget decisions live.

We will never tell you multi-touch is “too complex.” We will tell you precisely what each model can and cannot prove, and reconcile both against your backend where it's connected.

HOLD US TO IT In any QBR, point at a channel and ask what each model says about it. Expect two numbers and a reason.

MEASUREMENT & PROOF

Q09

How do you reconcile platform FTDs with my backend?

// THE SHORT ANSWER

Monthly, in writing, with every gap explained.

// IN PRACTICE

The reconciliation checks the boring things that create the gaps: event definitions, attribution windows, time zones, currencies, deduplication, naming and UTM logic — then compares platform-attributed FTDs against your backend-confirmed customers, line by line.

Until your backend is connected, platform FTDs are labelled platform-attributed in every report we send. Numbers that have never been checked don't get presented as facts.

HOLD US TO IT Ask for last month's reconciliation note. If a desk has never written one, their numbers have never been audited.

MEASUREMENT & PROOF

Q10

When did you last recommend cutting a channel that looked good?

// THE SHORT ANSWER

We cut channels that look good — when the path data shows the credit is borrowed. WE BRING: A REAL CHANNEL-CUT STORY

// IN PRACTICE

A channel earns a cut when it converts only on last-touch, when its players don't deposit twice, when frequency has passed the point of persuasion, or when pause tests show the conversions would have happened anyway. We run those pause tests and spend reviews as standing practice.

The specific story — which client, which channel, what the data showed, what happened after the cut — we tell in person. Our clients' accounts are real businesses, not anonymised marketing material. Ask, and you'll get names, numbers and the follow-up.

HOLD US TO IT

Ask the question in the room. A desk that has never recommended spending less is selling media, not managing it.

MEASUREMENT & PROOF

Q11

How would you design an incrementality test on my account?

// THE SHORT ANSWER

Designed to your volume — or we tell you you're not ready yet.

// IN PRACTICE

The design follows your market count, conversion volume and budget: geo splits and matched markets where you have several territories, audience holdouts where you don't, pause-reactivation tests where budget discipline demands it. Every design is written before launch: hypothesis, test and control, KPI, runtime, minimum volume, exclusions and the decision rule.

If your account can't yet support a statistically defensible test, the honest recommendation is a measurement-readiness phase and a clearly-labelled directional test — not a fake experiment that flatters everyone.

HOLD US TO IT

Ask what minimum volume the test needs. A desk with a real methodology answers with a number.

MEASUREMENT & PROOF

Q12

What’s your attribution window, and who can change it?

// THE SHORT ANSWER

It’s in the contract. Changing it requires your sign-off.

// IN PRACTICE

The window is agreed per objective before launch — tighter for lower-funnel registration and FTD activity, longer where CTV and DOOH influence is genuinely part of the journey — and it's documented where both sides can see it.

It does not move after results are visible. Window-shifting is the oldest trick in performance reporting, which is exactly why the change control belongs to you, not to us.

HOLD US TO IT — Check your current contract for the window. If it isn't written down, it's already movable.

MEASUREMENT & PROOF

Q13

Which of my KPIs — NRC, FTD, RDC — do you actually optimise toward?

// THE SHORT ANSWER

The one your economics point to — usually FTD, with RDC as the quality check.

// IN PRACTICE

Cheap registrations that never deposit are a cost, not a result, so we don't stop at them. The working hierarchy: qualified exposure up top, registration started and completed mid-funnel, FTD as the primary optimisation event, and repeat deposit as the standing check that the players we buy are worth having.

Where FTD or RDC signals aren't yet available in-platform, we build the plumbing — postbacks, offline conversion imports, CRM reconciliation — rather than optimising to whatever happened to be measurable.

HOLD US TO IT Ask which platform event the algorithm is literally optimising toward today. The answer should match your P&L.

MONEY, OPERATIONS & FIT

Q14

Break down every fee between my budget and the publisher.

// THE SHORT ANSWER

Line by line, on one page — shown before you sign.

WE BRING: A SAMPLE FEE BREAKDOWN

// IN PRACTICE

Media cost, platform fees, data costs, verification, ad serving where it applies, and ThumbAd's fee: itemised, with pass-through costs identified as pass-through. The structure varies by channel and buying method; the transparency doesn't.

You should be able to trace every euro from your budget to a publisher and account for the difference. If any line of that trace requires trust instead of arithmetic, we've failed the question.

HOLD US TO IT Ask for the fee sheet in the first commercial conversation, not the last.

Do you take rebates or arbitrage media?

// THE SHORT ANSWER

One disclosed model, agreed in writing — you always know exactly how we're paid.

// IN PRACTICE

We operate as a transparent managed service: our compensation is stated in the agreement, and whether media is bought at cost or under another agreed structure is written where you can point at it. No compensation that depends on you not reading the contract.

The test of any model is whether our incentive bends away from your efficiency. Ours is built not to: nothing at ThumbAd breaks when your CPA falls.

HOLD US TO IT Ask us to state, in one sentence in the contract, every way ThumbAd is compensated on your account.

Who actually trades my account, and how senior are they?

// THE SHORT ANSWER

Named people, introduced before you sign — and you can ask their account load.

// IN PRACTICE

ThumbAd is a senior-led desk: the people in your kickoff are the people in your account. You'll know by name who owns strategy, trading, reporting and escalation, and complex work — multi-market iGaming, pDOOH, attribution — gets senior hands, not junior handoffs.

Ask the account-load question directly and we will answer it: the current number of accounts carried by the specific traders who would be on yours, given in the meeting rather than printed here, because the honest figure moves as the book moves and a number in a brochure would be stale the week after we wrote it. Attention is the real product, and a desk that won't give you the number is telling you it's a bad one.

HOLD US TO IT Meet the actual trader before signing. “The team” is not a person.

MONEY, OPERATIONS & FIT

How do you handle responsible-gambling and licensing rules in my markets?

Q17

// THE SHORT ANSWER

We ask about your licences before we show you a single reach number.

// IN PRACTICE

Market context comes first: licensed territories, geo and age restrictions, responsible-gambling messaging, self-exclusion obligations, bonus and creative rules, platform and app-store policies. All of it shapes the plan before activation — it is planning input, not post-launch friction.

We do not learn on your licence. In this category, a compliance mistake doesn't cost budget; it costs the thing the whole business stands on.

HOLD US TO IT

Notice which comes first in our pitch: questions about your licence conditions, or slides about reach. The order is the tell.

MONEY, OPERATIONS & FIT

What does your weekly and quarterly reporting actually contain?

Q18

// THE SHORT ANSWER

Decisions, not metrics dumps. WE BRING: SAMPLE REPORTS

// IN PRACTICE

Every report answers four questions: what happened, why, what we did about it, and what happens next. Weekly: performance versus target by market, delivery and pacing, what changed in the account and what we changed in response. Monthly adds the full funnel, backend reconciliation and cohort maturation. Quarterly adds path-to-conversion findings, test verdicts, one honest miss, and the plan.

One consistent visual system across all of it, so a number means the same thing in the weekly note and the board deck.

HOLD US TO IT Ask for a real (anonymised) example of each cadence before signing. We'll send all three.

Why shouldn't I just take this in-house?

Q19

// THE SHORT ANSWER

Sometimes you should — and we'll tell you which parts.

// IN PRACTICE

In-housing works when you have the scale, the platform access, the technical support and — hardest of all — the sustained senior attention to run it well. When you do, the honest move is to say so, and we do: several capabilities belong inside a mature operator.

What tends to stay specialist: advanced setup and troubleshooting, market expansion, pDOOH and PMP development, deep Trade Desk expertise, attribution and path-to-conversion analysis, independent performance review, and training your team during the transition so no learning is lost. We'd rather be your specialist layer for years than your everything for one contract cycle.

HOLD US TO IT Ask us to map which of your functions should be internal in two years. Watch whether the answer costs us revenue.

MONEY, OPERATIONS & FIT

Walk me through an account you lost and why.

Q20

// THE SHORT ANSWER

Yes — a real one, with what we changed afterwards.

WE BRING: THE STORY, IN PERSON

// IN PRACTICE

Every desk that has traded for over a decade has lost accounts: to in-housing, to budget cuts, to market exits, to measurement disagreements, to procurement. Pretending otherwise is the least trustworthy sentence in any pitch.

We tell the story in person — a specific account, what happened, what we could have done differently, and the operating change it produced — because real losses involve real clients, and we don't anonymise them into content. What we'll commit to in print: the story exists, it's specific, and it ends with a process improvement you can still see in how we run accounts today.

A partner who never reports a miss isn't miss-free. They're hiding them — and the guide you're holding lists that as red flag number ten.

HOLD US TO IT Ask in the first meeting. The quality of this one answer predicts the honesty of every future report.

Now score us

The companion guide ends with a six-criteria scorecard and a 90-day evaluation plan. We're inviting you to run both — on us. Score each criterion 1–5 with evidence attached; anything scored on impression defaults to a 3. We consider 24 of 30 the bar worth scaling with, and we expect to be held to it.

Measurement honesty

Path exports on request, models compared openly, misses reported unprompted. Evidence: Q07, Q08, Q09, Q20.

Audience craft

Behavioural segments you own, tested against controls, documented. Evidence: Q01, Q02, Q03.

Market fluency

Licences, payment rails and market rules asked about before answered about. Evidence: Q04, Q17.

Fee transparency

Every euro accounted for, in writing, before signature. Evidence: Q14, Q15.

Communication quality

Named senior people; reports that explain decisions. Evidence: Q16, Q18.

Compounding evidence

Documented taxonomy, versioned learning, honest verdicts that persist. Evidence: Q03, Q05, Q11.

And run the same scorecard on your incumbent in parallel. Whatever you decide, ninety days from now you'll hold two evidence files instead of two pitches — which was the entire point of writing any of this down.

Twenty answers. Zero footnotes hiding.

Everything a desk hopes you won't ask, answered in writing before you asked it. The meeting version comes with exports, fee sheets and one uncomfortable story told properly.

The guide	The iGaming Operator's Guide to Programmatic — the 20 questions and how to score any partner
This document	Our answers, on the record, with the artifacts named
The meeting	Bring both. Score us live. Keep the evidence either way.

// BOOK THE EXAM

One hour. Your questions, our exports.

Independent programmatic trading desk, Oslo, since 2012. The Trade Desk across Europe, Africa, North America and Latin America — betting and gaming at the core.

thumbad.com