QUEENSLAB
ESTD Group 2025

THE COMMERCIAL MODEL

RETENTION AS SUPPLY RISK SERIES

DECISION MEMO

PART I.

AUGUST 2026

QUEENSLAB GROUP
Commercial Intelligence Series

DECISION MEMO — PART I OF II
RETENTION AS SUPPLY RISK: A TWO-PART SERIES

Part I — The Commercial Model

A two-segment model of plasma donor churn

Prepared for: Commercial / Sourcing Leadership, Plasma Collection Operators

Data basis: Proxy donor-level dataset, anchored to published industry benchmarks (illustrative model — not derived from any single operator's internal data)

Series: Part II — "The Floor Reality" — covers the operational failure modes and throughput bottleneck behind these numbers, for ops/center-manager audiences

The problem

Plasma collection economics assume retention is a single lever: keep donors happy, keep the liters flowing. In practice, donor populations mostly split into two behaviorally distinct groups that respond to opposite interventions — and treating them as one undifferentiated pool means misallocating retention spend on the segment least likely to respond to it.

This split isn't a plasma industry quirk. It mirrors a documented distinction in consumer behavior — transactional versus relational exchange (Dwyer, Schurr & Oh, 1987): some customers relate to a business as a series of discrete trades; others fold the relationship into their own identity. Two years on donor room floors made this distinction visible well before any literature confirmed it — donors describing "our" center rather than "the" center, treating a missed week as a small personal absence rather than a skipped transaction. That maps directly onto social identity theory (Tajfel & Turner, 1979): people who've folded a group into their self-concept talk about it collectively; people in a purely transactional relationship don't.

It also explains why a failure event doesn't hit every donor the same way. Psychology's well-documented finding that negative experiences weigh more heavily than equivalent positive ones — "bad is stronger than good" (Baumeister et al., 2001) — only bites for someone already invested in the relationship. A donor who's never framed the center as "theirs" loses nothing beyond the visit itself; a donor who experiences a failed venipuncture or a long wait as something closer to a small betrayal. That asymmetry was visible on the floor long before it showed up in the data as a 0%→81% swing in churn risk.

And it explains why incentive-chasing persists even as the national registry closes off its main advantage: variable-ratio reinforcement (Skinner) — intermittent, unpredictable rewards reinforce behavior more durably than steady ones. Donors chasing bonuses across centers described their own logic in exactly these terms, shopping for whichever center offered the better deal that month, independent of relationship. The registry constrains the mechanics of that strategy; it doesn't remove the psychological reward structure driving it. This blind spot compounds at the industry level: when donor systems only surface frequency, volume, and referral counts, every operator ends up reading the same incomplete instrument — which is a large part of why retention strategy across Czech operators looks so interchangeable.

The two segments

	Cash Optimizer	Identity-Loyal
Primary driver	Incentive differential (now capped by national registry)	Failure / experience events
Donation frequency	High	Moderate / stable
Vein-wear trajectory	Fast depreciation — self-limiting exit	Slow depreciation

Retention lever	Limited — registry already constrains upside	Failure prevention, staff continuity
Language marker	"I" — individual transaction framing	"We" — collective / belonging framing

Cash optimizers

Incentive-driven, cross-shop between centers for bonuses and referral payouts, treat donation as supplemental income rather than commitment. Historically able to maximize earnings by donating across multiple operators simultaneously; this upside is now structurally bounded by the Czech national plasma donor registry, which tracks cross-center donation frequency to enforce safety limits — a regulatory development that caps, without eliminating, this segment's arbitrage strategy.

Identity-loyal donors

Single-center, relationship-driven, self-identify with the center ("we" rather than "I" in how they describe it). Tolerant of minor friction; retention here is driven almost entirely by experience quality, not incentive size.

The split, in numbers

Applied to the 3,001-donor proxy population: 1,334 donors (44%) classify as Cash Optimizers, 907 (30%) as Identity-Loyal, and 760 (25%) as one-time/unclassified (a single donation, never establishing either pattern). Cash Optimizers generated 17,577 donations against Identity-Loyal's 5,772 — roughly three times the volume, at just under one and a half times the donor count. (Dollar/CZK figures are deliberately held for the Financial translation section below, where they're anchored to an external benchmark rather than the dataset's internal revenue assumption — see the note there.)

The failure layer

Four categories of collection failure act as the primary churn trigger for Identity-Loyal donors and, cumulatively, erode the Cash Optimizer segment's physical capacity to donate:

- Donor readiness — inadequate prep, travel restriction, temporary medical deferral
- Procedural — failed venipuncture / re-stick, separator malfunction
- Lab / QC — failed screening, hemolytic or lipemic sample rejection
- Adverse event — vasovagal reaction or collapse during / after draw

A fifth, chronic factor compounds these: cumulative donation count degrades vein quality over a donor's lifecycle, raising procedural difficulty and pain over time. Because Cash Optimizers donate at the highest frequency, they deplete this capacity fastest — meaning the segment most driven by incentive economics is also the one most likely to self-exit physiologically, independent of any competitor's offer. (This wear effect is a literature-supported hypothesis about plasmapheresis donor fatigue; it is not yet visible in this dataset's raw failure rates by visit sequence, and is flagged here as a direction for future validation rather than a current finding.)

The evidence: churn risk by failure exposure

A donor's Churn_Risk_Tier (Active/Regular/At-risk/Lapsed — see Methodology) responds very differently to a failure event depending on segment, at matched donation frequency (donation-frequency tiers run, in ascending order, rare → occasional → regular → super_active):

Frequency tier	Segment	No failure event	Had a failure event
regular	Identity-Loyal	0.0% high-risk	
regular	Cash Optimizer	17.2% high-risk	46.0% high-risk
occasional	Identity-Loyal	10.1% high-risk	
occasional	Cash Optimizer	69.3% high-risk	100% high-risk

At the same donation frequency, a failure event takes a regular-cadence Identity-Loyal donor from effectively zero churn risk to over 80% — roughly double the relative jump seen in the equivalent Cash Optimizer cohort. This is the quantitative basis for the recommendation below: failure prevention is a retention lever with a measurably larger effect on the Identity-Loyal segment.

The financial translation

One note on method: The only externally anchored benchmark used below is Jaworski, P., "America's Plasma Contribution to the World: 2025," a public estimate of plasma selling price at ~$190/liter (±$10) — the same source used in the Paradox and SDLA case studies — converted at ~21.2 CZK/USD (≈ 4,028 CZK/L) and applied to this dataset's own measured average donation volume. Volume itself (liters collected) is a direct, measured quantity in the dataset and needs no assumption.

Of the 1,334 Cash Optimizer donors, 297 (22.3%) sit in the At-risk/Lapsed tier. At this segment's median donation frequency and average yield per draw, that represents roughly 6,600 L/year at risk — about 26.6M CZK (~$1.25M) in projected annual plasma revenue, on the external benchmark. Of the 907 Identity-Loyal donors, 443 (48.8%, nearly half the segment) sit in At-risk/Lapsed, representing roughly 4,500 L/year at risk — about 18.2M CZK (~$860K). Identity-Loyal's absolute exposure is smaller, but at almost double the share of its own donor base, which is the segment's real vulnerability.

The registry what-if

This dataset predates the national registry and reflects single-center behavior only — it cannot directly observe cross-center arbitrage. To size the registry's likely effect, we model a sensitivity range: the share of Cash Optimizer donation volume plausibly enabled by pre-registry cross-center hopping beyond a single center's safe cadence.

Assumed arbitrage share	Volume reduction (L)	Value reduction (CZK / $, @190/L)
10%	1,372	5.53M CZK (~$261K)
20%	2,744	11.05M CZK (~$521K)
30%	4,116	16.58M CZK (~$782K)

Read this as the ceiling the registry places on Cash Optimizer segment value going forward, not a loss already incurred — the segment's post-registry volume settles somewhere below its pre-registry historical total (13,722 L across the segment), with the exact point depending on how much of that history depended on now-closed cross-center access. This range is an assumption-driven sensitivity, not a dataset-derived finding, and should be recalibrated against real cross-center donation data if it becomes available.

The recommendation

Retention spend should be split by segment, not be applied uniformly:

- Incentive / bonus spend — largely wasted on Identity-Loyal donors (already retained on relationship terms); increasingly capped in effectiveness on Cash Optimizers by registry-enforced frequency limits
- Experience / failure-prevention investment (venipuncture technique, wait-time reduction, staff continuity) — the higher-leverage lever for the Identity-Loyal segment, and the segment most worth protecting for long-run supply stability

Methodology

Segment (Cash Optimizer / Identity-Loyal) is a constructed field, not a directly observed one — the source data records donation frequency and referral usage, not stated motive. The construction rule: super_active frequency OR ≥ 1 referral → Cash Optimizer; regular/occasional frequency AND zero referrals → Identity-Loyal; single-donation ("rare") donors are held out as unclassified, since one visit is insufficient to express either pattern.

Churn_Risk_Score is likewise engineered, replacing the source dataset's original Status field (which we found was assigned purely from frequency tier at generation time and did not respond to failure history at all — an artifact that would have made any failure→churn claim circular). The new score blends donation recency and per-donor failure rate, weighted by segment to test the hypothesis directly: 60% failure-weighted / 40% recency for Identity-Loyal, 20% failure-weighted / 80% recency for Cash Optimizer. Scores are bucketed into quartiles (Active/Regular/At-risk/Lapsed) across the full population. This weighting was chosen to test the causal story in this memo, not discovered independently of it — stated plainly so the evidence table above is read as a model result, not a raw fact.

Model scope note

This model uses a synthetic proxy dataset constructed to reflect publicly reported industry deferral and adverse-event rate ranges. It is designed to demonstrate a reusable analytical framework, not to represent any single operator's actual donor population.

Continued in Part II — "The Floor Reality": the failure modes, the throughput bottleneck, and the booking-system fix behind the numbers above, written for ops and center-manager audiences.

QUEENSLAB
ESTD Group 2025

INTELLIGENCE REFINED
GROWTH REDEFINED

Created by QueensLab Group s.r.o.