



M A R G I N R R E S E A R C H

RESEARCH PAPER 2026-07 · THIRD EDITION

Discounted Margin Flow

Developing a new transparent model of margin health for dental practices, with a retrospective application to a three-site group

Paul-Matthieu Castro

Marginr Research

Third edition — August 2026

doi.org/10.5281/zenodo.21907533

This edition supersedes the July 2026 first and second editions; Appendix D records all changes between editions.

About this paper. Research Paper 2026-07 specifies the Discounted Margin Flow method and applies it to one retrospective engagement. The subject group is a Marginr client. The owner consented to anonymised use, and the case is anonymised. DMF is an operating diagnostic. It is not an appraisal, a quality-of-earnings review or a guarantee of recovery. Appendix C states the conflict, funding and data disclosures. Appendix D records every change between editions, and Appendix E publishes the follow-up record schema for the validation programme in Section 8.

© 2026 Paul-Matthieu Castro. All rights reserved except as licensed below.

Published under Creative Commons Attribution-NoDerivatives 4.0 International (CC BY-ND 4.0). <https://creativecommons.org/licenses/by-nd/4.0/>

Marginr™, Marginr Research™, Discounted Margin Flow™ and DMF™ are unregistered trade marks of Paul-Matthieu Castro. No trade mark rights are licensed.

Executive summary

What we found. This paper applies the Discounted Margin Flow method to one real, anonymised engagement: a three-site general dental group. The method read the group's own invoices, statements and contracts. It found **\$432,000 a year** of margin being lost or left unclaimed.

Why we sort it. One large number is easy to publish and impossible to defend. When a practice is sold, the buyer's accountants take that number apart line by line, in a room where the owner cannot answer. We do that sorting first, on the owner's side of the table. Every amount below is real money. The tiers describe what stands between the owner and it.

Tier 1 — Recoverable register: \$185,000 a year

Every item here is written on a document we read, repeats every year, and has a route the owner can start. The tier divides by one question: who has to agree?

- **Owner-only — \$40,000 a year.** Charges for things the group no longer uses, and a volume rebate already earned under terms already signed. The owner signs, and the money stops leaving. Assume the worst case: every supplier refuses to negotiate, every payer refuses to move, every habit reverts. This \$40,000 is still there. It depends on nobody else, and it is the floor of the engagement.
- **One agreement away — \$145,000 a year.** The same quality of evidence, one extra step. We can show the invoice where one site pays \$118 for the item another site pays \$154 for. The cheaper price is not a theory; it exists on paper, inside the group. What we cannot promise is that the supplier, the laboratory or the card processor agrees to match it. So this money is documented, repeating and identified, and it stays a candidate until somebody signs. We do not count it as recovered, and no honest report should.

Tier 2 — Conditional opportunity: \$247,000 a year

Also real, and worth knowing. Each item needs something outside the owner's control to change first: patients accepting different prices, a national insurer amending a contract it already signed, empty chairs actually filling, staff sustaining a new habit. One item, \$45,000 of unpaid old bills, is real money that can be collected only once, so it is not an annual rate at all. This tier is a shelf, not a reject pile. An owner who knows what sits on it can take an item down the day the condition changes: when a contract renews, when a lease ends, when a payer opens negotiations. What it must never be is added to Tier 1 and quoted as a single number.

\$432,000 / yr identified




Figure 1. The same \$432,000, sorted by what stands between the owner and the money. Nothing is discarded. Each tier states the condition attached to that amount.

What the sorting costs us. Of every \$100 found, about \$8 passes a buyer's quality-of-earnings review unchallenged: \$36,000, after a 10% realisation haircut on the owner-only tier, supporting \$180,000 to \$288,000 of illustrative enterprise value. We publish the \$8 and the route to it: a number that shrinks under examination was never worth having. Identified is not bankable.

On the 8%. That rate is a property of this practice, not of the method. This group's losses sat mainly where a counterparty has to agree. A practice whose losses sit in cancellable contracts and unclaimed entitlements converts a far higher share. We report the rate case by case, and we do not publish one figure as typical.

What this is not. Not an appraisal, and not a substitute for a quality-of-earnings review. No implementation outcomes exist yet: at the reporting date none of the five action gates was resolved. Section 8 publishes the validation programme that will produce outcome data.

Abstract

Three methods commonly price an independent dental practice: a percentage of collections, seller's discretionary earnings, and a multiple of adjusted EBITDA. Each one answers a transaction question. None of them decomposes the operating gap between realised and attainable margin. This paper specifies Discounted Margin Flow (DMF). DMF is an operating diagnostic with three distinct layers: a stock measure of margin capture, a register of document-supported recurring opportunities, and conditional scenarios for the present value and transaction sensitivity of implemented fixes.

The specification imposes two constraints. First, a partial expense extraction can identify dollar gaps, but it cannot produce a valid total-margin ratio unless omitted costs are restored. Second, a five-year persistence assumption does not make a saving permanent. In the retrospective case, the engagement identified approximately \$432,000 of annual opportunity, measured after the fee-schedule restatement in Section 5.3. Of that amount, \$185,000 qualifies for the candidate register, conditional on documentary action gates, and \$247,000 remains diagnostic context. Sorted by execution route, Tier 1 is the \$185,000 recoverable register, of which \$40,000 is owner-only and \$145,000 is one agreement away, and Tier 2 is the \$247,000 of conditional opportunity. At the displayed rounding level, the register supports a gross RMV scenario of approximately \$309,000 to \$423,000 across 25% to 8% discount rates, before unquantified switching costs. The contractual subset supports an illustrative EVF sensitivity of \$180,000 to \$288,000 at 5–8 times EBITDA. That sensitivity is conditional on implementation, persistence, buyer acceptance as normalised EBITDA and an unchanged multiple.

The paper uses reported EBITDA as a provisional full-margin proxy and holds omitted expense categories at their current level. On that basis the internal-frontier MCR is approximately 71%, a provisional upper bound pending complete P&L mapping. The paper carries no external stress case. The component bands available for one are not substantiated by the cited ADA survey workbook. The staff band among them is also not defined on the same basis as the labour line extracted from the engagement file (Section 5.7).

How to read it. An owner should read Sections 1 and 5 with the register tables. A deal-side reader should read Sections 4 to 6 with Appendix B; Section 5.6 gives per-line switching-cost breakevens and Section 6 gives a buyer-style realisation scenario.

1. Introduction

A transaction multiple asks what a buyer may pay for a practice relative to other practices. An owner between transactions asks a different question. Which margin losses are evidenced? Which of them can the owner act on? How much economic significance should attach to fixing them? A percentage of collections is margin-blind. Adjusted EBITDA responds to margin, but it does not identify the composition, actionability or persistence of the gap.

DMF answers the second question. It is an intrinsic-style diagnostic, because its primary benchmark comes from the subject's own observed performance and documents. It is not a full intrinsic valuation of the enterprise, and its optional transaction translation remains a relative-valuation sensitivity.

CLAIM DISCIPLINE

The paper separates observed inputs, calculated outputs, judgment assumptions and conditional scenarios. A document can support a price gap without proving that the full annual recovery will be realised.

The three layers carry different claim levels, and the paper never adds them together. The stock layer measures capture and reports a ratio. The flow layer records recurring opportunities that pass three document tests and an action gate. The scenario layer discounts the register under stated windows, haircuts and rates. Section 4 specifies each layer. Section 5 applies all three layers to the case. Sections 6 and 7 report the sensitivity of the result and the limits on it.

Two restatements were made after first publication, in the fee-schedule line and in the labour line. Both sit in lines that the specification had already excluded from the candidate register, the RMV scenarios and the EVF sensitivity, so neither one reaches a claim-bearing output. Section 5.3 states that containment property, and Appendix D records every change between editions.

2. Prior Methods and Their Limits

2.1 Percentage of collections

The collections rule of thumb requires one readily available input and remains useful as a quick market cross-check. Its limitation for operating diagnosis is structural: collections contain no information about the cost required to produce them. Two practices with the same collections and different overhead receive the same result.

2.2 Earnings multiples

Normalised EBITDA and SDE are more informative because they reflect profitability. Their correct use depends on the earnings basis, buyer type, owner dependence, transaction structure and quality-of-earnings adjustments. Advisor publications in 2026 report materially different bands by size and structure. They are directional market observations, not a coefficient that converts every prospective saving into certain value.

2.3 Capitalised earnings

Capitalised earnings is closer to an intrinsic approach because it relates value to an earnings stream and a capitalisation rate. A capitalisation rate is not interchangeable with a loan rate or a short-horizon discount rate. Lower rates produce higher present values; comparing an 8% discount rate with a 25–31% capitalisation range cannot, by itself, establish conservatism.

2.4 Relative and intrinsic questions

Relative valuation embeds assumptions about risk, growth, normalisation and comparability in the selected multiple. An explicit discounted model makes assumptions visible but does not make them accurate. DMF therefore uses discounted scenarios to order and scale documented opportunities, while preserving transaction multiples as a separate, conditional sensitivity.

Method	Primary base	Comparison	Margin diagnosis	Principal limitation here
% of collections	Revenue	Market rule	No	Ignores cost structure
SDE / EBITDA multiple	Normalised earnings	Transaction market	Level only	Does not decompose the gap
Capitalised earnings	Normalised earnings	Risk / cap rate	Level only	One rate does not identify mechanisms
DMF	Margin components	Self first	Gap and composition	Requires complete extraction and calibrated assumptions

Table 1. The methods answer different questions and should not be treated as substitutes.

TERMINOLOGY

This paper uses *operating diagnostic* for DMF, *scenario present value* for RMV, and *transaction sensitivity* for EVF. It does not claim to appraise the practice.

3. Theoretical Foundations

3.1 Visibility of assumptions

The intrinsic-valuation tradition, commonly associated with Williams (1938) and surveyed by Damodaran (2005), values an asset from its expected cash flows at a risk-adjusted rate. DMF borrows the discipline rather than the full object: it discounts a defined set of potential margin recoveries, not the enterprise's total cash flows. Its methodological benefit is auditability. Every duration, haircut and discount-rate scenario has a named position in the calculation.

3.2 Frontier efficiency

Farrell (1957) and Charnes, Cooper and Rhodes (1978) established the idea of measuring performance relative to an empirical frontier. DMF is inspired by that tradition but does not estimate a DEA or stochastic-frontier model. It uses accounting identities and observed reference performance. The distinction matters: no efficiency score should be represented as a DEA result.

3.3 Dentistry evidence

Lipscomb and Douglass (1986) studied dental-production efficiency, and Chen and Ray (2013) applied DEA and stochastic frontier analysis to Colorado practice data. These papers establish that scale and cost efficiency can be studied in dental practices; they do not validate DMF's particular benchmark, windows or haircuts.

3.4 A group as a partially controlled comparison

A multi-site group provides evidence unavailable to a solo practice: the same owner may observe different prices, collection performance and cost ratios across sites. Shared ownership and some common protocols reduce, but do not eliminate, confounding. Case mix, provider mix, catchment, lease vintage, labour market, capacity utilisation and accounting allocation can still differ.

The stock layer uses a reference-site frontier: where sites are sufficiently comparable, one observed site supplies the component mix. It does not construct a synthetic practice by selecting the best value from a different site for every component. Item-level price comparisons may still use the lowest verified price for an identical item because the unit, vendor and specification can be matched directly.

SCOPE OF THE FRONTIER CLAIM

An observed internal benchmark increases evidential relevance; it does not prove transferability, optimality or causation. A ratio gap becomes a diagnostic prompt. A register claim requires item-level documentation and an action gate.

3.5 Reference selection

The primary hierarchy is internal reference site, then sustained own history, then an explicitly sourced external comparator. The levels are not mechanically mixed. External bands are reported separately as stress cases, because a more demanding band is not necessarily more comparable evidence. The case carries no external stress scenario, on the grounds set out in Section 5.7. The hierarchy itself is unchanged, and a properly sourced and definitionally matched external comparator would occupy rank 3 as specified.

4. Model Specification

4.1 Outputs and claim status

DMF is a diligence map, not a value bridge: it identifies where to look, not what to pay. Each output below carries its own unit and its own claim level. The case reports the identified total in two tiers: Tier 1, the recoverable register, which divides into an owner-only sub-tier and a tier that is one agreement away; and Tier 2, conditional opportunity, which waits on a condition the owner does not control.

Output	Definition	Unit	Status
Candidate register	Document-supported, owner-actionable recurring opportunity, subject to the action gate	\$/year	Case claim; underlying documents not published
RMV scenario	Gross discounted value under stated windows, haircuts and rates; net RMV additionally deducts switching costs	\$ present	Conditional scenario, not forecast
MCR	Full realised operating margin divided by full attainable operating margin	%	Requires complete or explicitly restored cost base
Dollar gap	Attainable margin less realised margin for the same component set	\$/year	May be reported on partial extraction
EVF sensitivity	Conditional enterprise-value effect if savings become accepted normalised EBITDA	\$ range	Illustration, not appraisal

Table 2. Outputs use different units and claim levels and are never summed.

4.2 Full stock layer

Let C be realised collections and X_c each disjoint operating-expense component expressed in dollars. Full realised operating margin is:

$$M = C - \sum_c X_c$$

Let A_{adj} be adjusted production, n^* the attainable net collection rate, x_c^* the attainable expense rate and B_c the documented base appropriate to component c . In this case, gross production is the base for laboratory and non-dentist staff ratios. Attainable margin is:

$$M^* = A_{adj} n^* - \sum_c (x_c^* B_c) \quad \text{MCR} = M / M^*$$

MCR is reported only when numerator and denominator cover the same complete expense partition, or when omitted costs are restored under an explicit assumption. Analytical output is not clamped: values above 100% and negative values carry information and are flagged.

4.3 Partial extraction

For a documented subset S , the partial margin M_S and partial attainable margin M_S^* can identify the dollar gap $M_S^* - M_S$. Their ratio is not a total MCR. If full realised margin M is available, current omitted costs are $K = M_S - M$. Assuming omitted costs have no attainable improvement gives:

$$M_{UB}^* = M_S^* - K \quad \text{MCR}_{UB} = M / M_{UB}^*$$

This is an upper bound on capture if omitted categories could contain additional savings. It is provisional when full realised margin is itself approximated by reported EBITDA.

4.4 Attainable hierarchy

Rank	Source	Definition	Use
1	Comparable reference site	One observed site selected under a documented comparability rule	Primary internal MCR
2	Own historical frontier	Sustained historical period; duration is a judgment parameter until calibrated	Primary when no comparable site exists
3	External comparator	Precisely cited dataset with matching definitions and population	Separate stress case, not mixed into ranks 1–2

Table 3. Confidence order and ambition are separate dimensions.

4.5 Flow layer: the candidate register

A finding enters the candidate register only when it passes three tests:

- **Document-supported.** The price, volume and recurring basis are traceable to identified source pages. The annual recovery may still be an estimate and is labelled accordingly.
- **Owner-actionable.** The owner can execute the required action without further discretionary approval. Accepting a valid written offer or exercising an existing cancellation right can qualify; asking a counterparty to negotiate does not.
- **Recurring.** The amount is an annual run-rate rather than a one-time balance.

A fourth action gate records prerequisites: quote validity, capacity, product equivalence, switching cost, contract term, notice period and any operational dependency. If a prerequisite is unresolved, the item remains a candidate and cannot be described as guaranteed.

Items that fail are retained as diagnostic context. The context taxonomy distinguishes:

Class	Meaning	Permitted language
Contractual exposure	Observed economic amount requiring counterparty agreement	Documented exposure
Modelled opportunity	Counterfactual based on demand, ratios or process assumptions	Opportunity estimate, not cash loss
One-time recovery	Existing balance that does not recur annually	Balance recovery

COMMERCIAL BOUNDARY

This paper does not define or create contractual rights. Any commercial guarantee must state its terms in the client agreement and cannot be inferred from a research table.

4.6 Recoverable Margin Value scenarios

For finding i , let L_i be the annual candidate amount, h_i an explicit realisation haircut, W_i a credit window capped by the five-year model horizon, r a hurdle-rate scenario and S_i the present value of incremental switching and implementation costs:

$$\text{RMV}_{\text{gross}}(r) = \sum_i L_i h_i \sum_{t=1 \dots \min(W_i, 5)} (1 + r)^{-t}$$

$$\text{RMV}_{\text{net}}(r) = \text{RMV}_{\text{gross}}(r) - \sum_i S_i$$

The case file does not quantify termination fees, onboarding, training, equipment, dual-running time or internal implementation labour. The case calculations are therefore labelled gross RMV; net RMV is not estimable from the available evidence.

Haircut and window are assigned independently. The haircut reflects realisation evidence and execution prerequisites; the window reflects persistence mechanism. Both remain judgment assumptions until replaced by observed outcomes.

Axis	Level	Assignment rule	Parameter
Haircut	High evidence	Amount and implementation route documented; only administrative friction remains	$h = 0.90$
Haircut	Moderate evidence	Amount supported, but switching, quality, capacity or vendor execution remains unresolved	$h = 0.70$
Haircut	Model-dependent	Realisation depends materially on behaviour, demand or counterparty discretion	$h = 0.40$
Window	Contractual	Existing right, accepted offer or completed cancellation is expected to persist	5 years
Window	Repriced	Implemented external price can be revisited or switching conditions can change	3 years
Window	Behavioural	Recovery requires repeated staff or owner action	1 year

Discount-rate policy

The relevant rate is a risk-adjusted hurdle selected for the cash-flow scenario, not automatically the practice's debt cost. When the engagement file lacks a documented hurdle rate, no single RMV is designated as the estimate. The case reports a range across stated scenarios. The 8% column is a lower-rate scenario and therefore an upper RMV scenario, not a conservative floor.

4.7 Exit Value of Fixes sensitivity

EVF is reported only as a conditional transaction sensitivity for implemented recurring savings accepted as normalised EBITDA:

$$\text{EVF} = [\sum L_i h_i] \times [m_{\text{lo}}, m_{\text{hi}}]$$

It applies only if the fix is implemented, evidenced in trailing or accepted run-rate earnings, expected to persist beyond the transaction date, accepted by the buyer as normalised EBITDA, and priced at the same multiple. Reaching year five in a scenario does not, by itself, establish permanence.

5. Retrospective Application

5.1 Engagement and evidence boundary

The subject is a general dental group with three sites, six dentists and approximately 35 full-time equivalents in one metropolitan market. Annualised gross production was approximately \$6.4 million, adjusted production approximately \$5.16 million, net collections approximately \$4.85 million and reported EBITDA approximately \$0.92 million. Site C operates with less standardised purchasing and administrative routines than Sites A and B.

The retrospective engagement recorded approximately \$540,000 of annual opportunity, restated to approximately \$432,000 on the fee-schedule basis set out in Section 5.3. This paper classifies that record under the DMF specification. The client documents are confidential and are not reproduced. The author performed both the engagement and the retrospective analysis. The case therefore demonstrates internal reconciliation and model behaviour, not independent replication of source evidence.

Implementation status	Realised amount	Persistence	Outcome evidence
Not observed	Not measured	Not measured	No post-engagement implementation ledger, follow-up invoices or outcome P&L was available

Outcome boundary. “Retrospective” refers to the timing of the source-document review, not to validated implementation outcomes.

5.2 Reclassification at displayed rounding

Identified annual opportunity	Candidate register	Diagnostic context
\$432k	\$185k	\$247k

The displayed candidate-register and context figures sum to the identified total. Because all are rounded, this is a reconciliation at presentation precision rather than literal dollar-exact proof. The classification does not establish that every opportunity amount was economically realised or recoverable.

5.3 Fee-schedule opportunity and percentile restatement

The fee-schedule line is reported at approximately **\$52,000** on a median basis. The first edition reported it at \$160,000 and described it as an uplift “to market median”. Those two statements do not agree. The engagement's fee table measured each code's gap against the regional **75th** percentile and summed those gaps to approximately \$227,700 of additional annual production. A PPO-cap adjustment then produced the \$160,000 figure. Recomputed against the **50th** percentile on the same code set, the additional production is approximately \$73,000, and the same PPO-cap ratio gives approximately \$52,000. The procedure volumes used in the recomputation are reverse-engineered from the original table's annual-impact column rather than independently observed. The restated figure is therefore an estimate, and the paper labels it as such.

Independent of the arithmetic, the line fails the candidate-register test in either version, because recovery depends on patient response, payer rules and price realisation. It is a modelled opportunity, not an observed cash loss. A narrower contract-shaped opportunity may exist where posted fees are below a payer's allowed amount. The engagement file did not isolate that amount. Under the whole-line rule, none is promoted without targeted re-extraction.

WHY THE RESTATEMENT DOES NOT PROPAGATE

The restatement changes the identified total and the context total. It leaves the candidate register, the RMV scenarios and the EVF sensitivity unchanged, because the fee line never entered any of them. An error confined to a context line cannot reach a claim-bearing output. That containment is the intended behaviour of the partition, and it is the strongest available argument for separating claim levels before quantifying them.

5.4 Candidate register

Finding	\$/yr	Shape	<i>h</i>	Evidence recorded	Action gate	Gate status
Cross-site supply price variance, 31 SKUs	35,000	Repriced	0.90	Same-item distributor invoices	Cheaper route must be available without new negotiation	Unresolved
Group volume-rebate tier	10,000	Contractual	0.90	Combined spend and tier terms	Eligibility must be automatic under written terms	Unresolved
Laboratory per-unit repricing	95,000	Repriced	0.70	Invoices and reference-lab quote	Quote validity, capacity, quality and switching cost	Unresolved
Cancellation of four recurring items	30,000	Contractual	0.90	Contracts and payment statements	Notice periods and termination rights	Unresolved
Merchant processing change	15,000	Repriced	0.70	Statement and alternative terms	Written offer, switching cost and implementation conditions	Unresolved
Candidate total	185,000	Conditional on all action gates; none resolved at the reporting date				

Table 4. A source record can support candidacy; unresolved gates prevent guarantee language.

5.4.1 Register composition by execution route (Tier 1)

At the reporting date, 0 of the 5 action gates was resolved. \$185,000 is the ceiling of a conditional register, not a pipeline of confirmed savings. The register also splits by execution route, and the two routes carry different certainty.

Execution route	\$/yr	Items	What the owner must do
Owner-only tier	40,000	Group volume-rebate tier \$10,000; cancellation of four recurring items \$30,000	Sign alone, under an existing written right
One agreement away	145,000	Supply variance \$35,000; laboratory repricing \$95,000; merchant processing \$15,000	Obtain a counterparty's agreement, or a valid written offer, first
Candidate register	185,000	Ceiling of a conditional register; 0 of 5 gates resolved at the reporting date	

Table 4A. The execution route, not the dollar amount, carries most of the register's risk.

The three one-agreement-away items were recorded in the engagement with execution routes that begin by asking a counterparty to negotiate. The owner-actionable test in Section 4.5 does not accept that route. Promotion of those items requires either a valid written alternative offer or an exercisable existing right.

5.5 Diagnostic context (Tier 2 — conditional opportunity)

This tier is a shelf rather than a reject pile: each item is real, and each waits on a condition outside the owner's control. It is reported separately so that it is never added to the register and quoted as one figure.

Item	\$/yr	Class	Reason excluded
Fee-schedule uplift to market median (<i>median basis, 5.3</i>)	52,000	Modelled opportunity	Demand and price-realisation contingent
Largest-payer write-off spread	75,000	Contractual exposure	Requires payer agreement
Aged accounts-receivable backlog	45,000	One-time recovery	Balance, not recurring annual rate
Chair utilisation and labour efficiency	40,000	Modelled opportunity	Ratio counterfactual; causation not established
Facilities and service-contract rebids	15,000	Modelled opportunity	Counterparty outcome not evidenced
Cross-site workflow replication	10,000	Modelled opportunity	Metric gap rather than documented recovery
Standing patient-receivables process	10,000	Modelled opportunity	Behavioural and not independently documented
Context total	247,000	Excluded from register and RMV scenarios	

Table 5. Context is economically relevant but not equivalent to recoverable cash.

Scope note to Table 5. The engagement record also referred to approximately \$750,000 of total payer contractual adjustments. That quantity is distinct from the displayed \$75,000 largest-payer spread and is not added to the partition. The distinction is stated to prevent a tenfold ambiguity.

5.6 RMV scenario computation

The engagement file contains no documented hurdle rate and no quantified switching-cost schedule. The table therefore shows gross RMV mechanics at 8% as a lower-rate scenario, while the claim-bearing result is the full scenario range in Table 7.

Finding	L	W	h	$A(W, 8\%)$	Gross RMV at 8%
Supply variance	35,000	3	0.90	2.5771	81,179
Rebate tier	10,000	5	0.90	3.9927	35,934
Laboratory repricing	95,000	3	0.70	2.5771	171,377
Cancellations	30,000	5	0.90	3.9927	107,803
Merchant processing	15,000	3	0.70	2.5771	27,060
Total	185,000	—	—	—	423,353

Table 6. Display-level reconciliation of the 8% scenario.




Figure 2. Gross RMV components at 8%. Green denotes $h = 0.90$; navy denotes $h = 0.70$.

Rate scenario	8%	12%	15%	25% stress
Gross candidate RMV	\$423,353	\$390,371	\$368,408	\$308,606

Table 7. With no documented hurdle rate or switching-cost schedule, the case reports gross scenarios rather than a net estimate.

SCENARIO INTERPRETATION

Across all displayed rate scenarios, gross candidate RMV is approximately \$309,000–\$423,000. Net RMV is lower by the present value of switching and implementation costs, and the case file does not support an estimate of it. The restatements described in Sections 5.3 and 5.7 do not touch any register input, so these figures stand as first computed.

5.6.1 Net RMV and breakeven switching costs

Net RMV is gross RMV less the present value of switching costs. Switching costs include termination fees, onboarding, training, dual-running and internal labour. The case file does not observe any of them. The paper therefore reports gross figures, and it reports a breakeven for each line instead of guessing a cost.

The breakeven S^* is the one-time cost today that makes the line's value zero. It equals the line's gross RMV. A reader with a better estimate of any switching cost can apply it directly: a line stays value-positive while its one-time cost stays below its breakeven.

Finding	Gross RMV at 8% (= breakeven S^*)	$S^* \div \text{annual savings } L$	Breakeven at 25% stress
Supply variance	81,179	2.32×	61,488
Rebate tier	35,934	3.59×	24,204
Laboratory repricing	171,377	1.80×	129,808
Cancellations	107,803	3.59×	72,611
Merchant processing	27,060	1.80×	20,496

Table 6A. Breakevens use the same L , h and W as Table 6. The stress column uses $A(3, 25\%) = 1.9520$ and $A(5, 25\%) = 2.68928$.

The ratio column states each breakeven as a multiple of the line's annual savings. A contractual five-year line tolerates a one-time cost of about 3.6 times its annual saving. A three-year line at $h = 0.70$ tolerates about 1.8 times. The stress column repeats the test at a 25% rate, where every line tolerates less. Switching costs act on the register in the same direction as the discount rate, and neither one is observed in this case.

5.7 Stock-layer results

The engagement documented six cross-site metrics. Site A was best on all six displayed metrics, so the case does not require a synthetic component-by-component frontier.

Metric	Site A	Site B	Site C	Reference
Net collection rate	97%	94%	89%	A
Laboratory, % gross production	11%	12%	16%	A
Supplies, % gross production	5.2%	6.1%	7.4%	A
Total labour incl. dentist compensation, % gross production	29%	32%	33%	A
Days in accounts receivable	41	49	63	A
Claim denial rate	4%	7%	11%	A

Table 8. Six metrics from the engagement record. The table carries no external band column; the paragraph below states why.

The labour line in Table 8 includes dentist compensation. The extracted line is \$1,984,000, or 31% of gross production, and the engagement report displayed dentists as 47% of the same labour pool. In standard dental benchmarking, a 25–30% band measures staff excluding dentists. Applying that band to an inclusive line is definitionally incompatible: six dentists and all remaining staff cannot be paid from 25% of production. The residual arithmetic confirms that the line is inclusive. Omitted current costs of approximately \$1.146 million could not carry six dentists' compensation together with all supply, occupancy, marketing and equipment cost. The paper therefore carries no external band comparison for this line and no stress case built on one. The approximately 49% stress MCR and the \$942,000 stress gap reported in the first edition are withdrawn and should not be used. The internal reference result is unaffected, because Site A and the group are measured on the same definition, so the comparison holds whatever the line includes.

The documented subset used for the MCR display contains collections realisation, laboratory and total labour. It produces partial realised margin of \$2.066 million, internal-reference partial attainable margin of \$2.445 million and a \$379,000 component gap. Supplies were not included in the dollar extraction, even though ratios were displayed. The paper therefore treats the three-component result as incomplete.

The paper uses reported EBITDA of approximately \$0.920 million provisionally as full realised margin. Current omitted costs are therefore approximately \$1.146 million. Holding those costs unchanged gives an internal MCR of approximately 71%. That figure is a provisional upper bound, pending complete P&L mapping.

Basis	Partial M^*	Less omitted K	Full M^* upper-bound basis	MCR upper bound	Dollar gap
Internal reference site	\$2.445M	\$1.146M	\$1.299M	~71%	\$379k

Table 9. Display-level calculation. The ratio is a provisional upper bound under unchanged omitted costs. The table carries no external band row.

INTERPRETATION OF THE STOCK LAYER

Restoring omitted costs changes the ratio while leaving the documented dollar gap unchanged, because the same cost term cancels in a difference but not in a ratio. A full P&L extraction is required for a final MCR estimate.

5.7.1 Bridge between stock gaps and flow items

The stock gap and the candidate register are two views of some of the same mechanisms. The bridge below reconciles them. The paper never adds the two together.

Stock component	Stock gap	Nearest flow item	Treatment
Collections	\$155,000	Payer write-off spread, \$75,000	Context; not promoted
Laboratory	\$96,000	Laboratory repricing, \$95,000	Same mechanism in two units; never added
Labour	\$128,000	Chair utilisation context, \$40,000	Ratio counterfactual; not a cash flow
Total	\$379,000	—	Diagnostic only; adding it to the register would double-count

Table 9A. The bridge makes the overlap visible and prevents double counting.

5.8 Conditional transaction sensitivity

The contractual-shape candidate items are the rebate tier and cancellations. Their combined annual amount is \$40,000, or \$36,000 after the stated 0.90 haircut. Merchant processing is classified as repriced because provider terms and switching conditions can change.

Advisor-published 2026 ranges vary by earnings basis and scale. The case reports approximately \$0.92 million of EBITDA, but the paper does not establish that it is fully normalised or that the group qualifies for a specific institutional buyer band. A 5–8 times range is therefore retained only as an illustrative sensitivity:

Haircut annual run-rate	At 5 times	At 8 times
\$36.0k	\$180.0k	\$288.0k

This sensitivity is conditional on four facts that multiplication alone does not establish: implementation, persistence, buyer acceptance as normalised EBITDA, and an unchanged multiple. If any one of them fails, the transaction effect is lower or zero. The paper never sums EVF with RMV or with register dollars.

6. Sensitivity and Model Risk

6.1 Discount rate

Across the full 8–25% display, gross RMV falls from \$423,353 to \$308,606. Short horizons make the calculation less rate-sensitive than a perpetuity, but the rate remains a judgment input when the required return is undocumented.

6.2 Windows, with the horizon fixed

Window case at 8%	Repriced window	Contractual window	RMV	Change
Minus one year	2	4	\$312,721	–26.1%
Specified	3	5	\$423,353	—
Plus one year, capped at five	4	5	\$503,103	+18.8%

Table 10. The upper scenario is capped at the model's five-year horizon.

6.3 Haircut sensitivity

A local plus-or-minus 0.10 change in each haircut, bounded at 1.00, changes gross 8% RMV to approximately \$370,014 and \$476,691. These scenarios remain gross of unquantified switching costs.

A buyer-style realisation scenario

External review from an M&A advisory practitioner (2026) indicates buyers commonly apply realisation discounts of 30–50%, not 10%. The model's own haircuts are judgment values, as Section 7.6 states. The row below shows the result under a buyer-conservative alternative: a uniform $h = 0.50$ on every register item, at 8%, with the specified windows unchanged.

Finding	L	W	$A(W, 8\%)$	Gross RMV at $h = 0.50$
Supply variance	35,000	3	2.5771	45,099
Rebate tier	10,000	5	3.9927	19,964
Laboratory repricing	95,000	3	2.5771	122,412
Cancellations	30,000	5	3.9927	59,891
Merchant processing	15,000	3	2.5771	19,328
Total	185,000	—	—	266,694

Table 11. Uniform 50% realisation at 8%. Gross RMV is \$266,694, about 37% below the specified-haircut scenario of \$423,353.

The same discount applied to year-one amounts gives \$92,500 on the register and \$20,000 on the owner-only tier. This scenario replaces the model's haircuts; it does not correct them. It is reported so that a deal-side reader can see the register at a realisation rate drawn from buyer practice rather than from the paper's own rules.

SCOPE OF THE RESTATEMENTS

The fee restatement and the removal of the external band comparison affect the identified total, the context total and the stock-layer external comparison. They do not touch L , h or W for any register item. Every RMV and EVF figure in this paper therefore stands as first computed.

7. Limitations

7.1 One retrospective case

The paper demonstrates the specified logic on one engagement performed and analysed retrospectively by the model's author. No post-engagement implementation or realised-cash outcomes were available. The case cannot establish predictive validity, external validity, causal recovery or typical register-to-context proportions.

7.2 Confidential evidence and reproducibility

The source invoices, contracts and extraction code are not published. The displayed calculations can be independently recomputed, but the documentary classifications cannot. Automated schema and build checks are process controls, not published evidence, unless the code and an auditable data representation are released or independently attested.

7.3 MCR remains provisional

The paper uses reported EBITDA as a proxy for realised margin, and the case file does not reconcile the two. The proxy may not share the exact definition of operating margin in Section 4. A full component P&L is required. The displayed 71% result is a provisional upper bound, pending complete P&L mapping, and it is not a final MCR estimate.

7.4 The paper carries no external component bands

The 2025 results workbook from the ADA Health Policy Institute's 2026 Survey of Dental Practice reports aggregate expenses and other practice statistics. It does not provide the 7–11% laboratory, 25–30% staff, 4–6% supply and 98–100% collection-rate bands used in the first edition. Those values were therefore never ADA-sourced. A second and independent reason removes them from the paper entirely. The staff band among them is not defined on the same basis as the corresponding line extracted from the engagement file, as Section 5.7 sets out. An unsourced band can still be labelled a judgment and kept as a stress case. A band that does not measure the same quantity as the subject line cannot be kept, because the comparison is not arithmetically meaningful at any confidence level.

7.5 Internal references remain confounded

Shared ownership reduces some differences but does not equalise case mix, provider mix, catchment, lease terms, labour markets or local operating maturity. Reference-site use requires a comparability screen and should be supplemented by item-level evidence. One specific confound is worth naming: laboratory ratios are sensitive to aligner volume, whose laboratory cost runs at a materially different share of case fee than conventional restorative work. Where aligner mix differs across sites, the laboratory ratio comparison is confounded, and the register's laboratory item should rest on the item-level invoice and reference-quote evidence rather than on the ratio.

7.6 Windows and haircuts are uncalibrated

Section 4.6 defines the assignment rules for the 5/3/1-year windows and 0.90/0.70/0.40 haircuts, but the parameter values remain judgments. Transparent rules make them contestable, not empirically correct. RMV remains a scenario until longitudinal evidence supports replacement parameters.

7.7 Transaction ranges are directional

Published dental multiples are sensitive to SDE versus adjusted EBITDA, owner dependence, buyer type, payer mix, earnings quality, deal structure and market date. EVF cannot substitute for an appraisal or quality-of-earnings review.

7.8 Re-identification and consent

The group consented to anonymised use, but the combination of site count, staffing and production may still carry re-identification risk. This edition removes further identifying detail from Section 5.1. Future case releases should document consent scope, rounding rules and a re-identification review.

7.9 Switching costs are not observed

The case does not quantify termination, onboarding, training, equipment, dual-running or internal implementation costs. Displayed RMV values are gross. Net RMV could be materially lower, and the paper does not estimate it. Section 5.6.1 states the breakeven one-time cost for each register line, so a reader who holds a better estimate can apply it without re-running the model.

7.10 Internal review found two errors after first publication

The author found both errors by re-reading the engagement report against the paper. No external check produced them, and both survived the first edition's own review. That record is evidence about the limits of single-author review, and it is the reason the programme below treats independent review as a requirement rather than an option. Appendix D records what changed in each edition.

8. Calibration and Validation Plan

Calibration should replace judgment parameters only after the outcome definitions, observation windows and evidence standards are fixed in advance. The programme below states the question, the minimum record that answers it, the timing of collection and the output it produces.

Question	Minimum record	Timing	Output
Implementation: was the action taken?	Dated evidence, action-gate status and implementation cost	30 and 90 days	Implementation rate by shape
Realisation: how much was recovered?	Pre/post unit price, or cash receipt on comparable volume	90 days and 12 months	Realised-to-claimed ratio with an uncertainty interval
Persistence: how long did it last?	Repeated measurement, with the reason for any failure	12, 24, 36 and 60 months	Survival by mechanism
MCR prediction: did the ratio predict improvement?	Full P&L mapping and the follow-up outcome	Annual	Out-of-sample validation, not in-sample fit

Table 12. The instrument and the timing are fixed before collection begins.

Four commitments govern the programme. First, development cases and validation cases stay separate, so no case both sets a parameter and tests it. Second, the follow-up record schema is published as Appendix E before data collection begins; it is the preregistered data dictionary. Third, items still active at measurement are treated as right-censored, and are neither assumed persistent nor counted as failed. Fourth, an independent reviewer checks the source classifications, sampling at least 20% of register items, before any outcome is published.

A 24-month follow-up cannot validate a five-year contractual window. Parameter updates should state the sample size, vintage, missingness and confidence interval. The programme also retains failed and unimplemented findings. A synthetic public case can make the complete calculation reproducible without exposing client information.

STATUS OF THE METHOD

Until this programme reports, DMF is a specified method with one retrospective application, not a validated tool.

9. Conclusion

DMF is a transparent operating-diagnostic framework. Its central contribution is the separation of observed dollar gaps, action-gated recurring candidates, contextual opportunities and explicit scenario assumptions. The framework is an operating diagnostic; it is not a quality-of-earnings substitute and its outputs are not appraisal inputs. Its claims stay bounded by the evidence available. The share of identified margin that survives to the defensible tier is a property of the practice examined, not of the method, and it is reported case by case.

In the case, \$185,000 of the restated \$432,000 identified total qualifies for the candidate register, conditional on documentary action gates. None of those gates was resolved at the reporting date. \$247,000 is retained as context. Under the stated windows and haircuts, gross candidate RMV is approximately \$309,000–\$423,000 across 25%–8% rate scenarios, before unquantified switching costs. The internal MCR is approximately 71%, a provisional upper bound, with a \$379,000 observed-component gap. The paper carries no external stress case.

No implementation outcomes were available for this case. The results show model behaviour, not realised recovery. The model's worth will depend on whether prospective follow-up shows that action-gated findings are implemented, recovered and sustained at rates consistent with the disclosed scenarios.

RESEARCH CONCLUSION

DMF's value lies in disciplined decomposition and transparent scenario design. Whether its candidate register predicts realised outcomes remains an empirical question for prospective validation. The two restatements offer one piece of indirect evidence for the design: both landed in lines the partition had already excluded from every claim-bearing output, and neither reached a result.

References

- American Dental Association, Health Policy Institute. (2026). *Income, Gross Billings, Expenses, Characteristics, Selected Staff Wages: Selected 2025 Results from the Survey of Dental Practice* [Excel workbook]. <https://www.ada.org/resources/research/health-policy-institute/dental-practice-research>. Accessed 28 July 2026.
- Charnes, A., Cooper, W. W., & Rhodes, E. (1978). Measuring the efficiency of decision making units. *European Journal of Operational Research*, 2(6), 429–444. [https://doi.org/10.1016/0377-2217\(78\)90138-8](https://doi.org/10.1016/0377-2217(78)90138-8).
- Chen, L., & Ray, S. C. (2013). Cost efficiency and scale economies in general dental practices in the US: A non-parametric and parametric analysis of Colorado data. *Journal of the Operational Research Society*, 64(5), 762–774.
- CT Acquisitions. (2026). *Dental and DSO M&A Multiples Report 2026: Valuation Benchmarks by Size and Segment*. Updated 29 June 2026. <https://ctacquisitions.com/guides/dental-dso-ma-multiples-2026/>. Accessed 28 July 2026.
- Damodaran, A. (2005). Valuation approaches and metrics: A survey of the theory and evidence. *Foundations and Trends in Finance*, 1(8). Working-paper version dated November 2006: <https://pages.stern.nyu.edu/~adamodar/pdfiles/papers/valuesurvey.pdf>.
- DDSMATCH South. (2025). *Dental Practice Valuation Explained for Practice Owners*. 22 April 2025. <https://ddsmatchsouth.com/dental-practice-valuation-explained/>. Accessed 28 July 2026.
- Farrell, M. J. (1957). The measurement of productive efficiency. *Journal of the Royal Statistical Society: Series A*, 120(3), 253–281. <https://doi.org/10.2307/2343100>.
- FOCUS Investment Banking. (2026). *Dental Practice EBITDA Multiples 2026*. 27 April 2026. <https://focusbankers.com/dental-practice-ebitda-2026/>. Accessed 28 July 2026.
- Lakeshore Dental Brokers. (2026). *How to Value a Dental Practice: Key Metrics and Multiples*. <https://www.lakeshoredentalbrokers.com/resources/how-to-value-a-dental-practice/>. Accessed 28 July 2026.
- Lipscomb, J., & Douglass, C. W. (1986). Are larger dental practices more efficient? An analysis of dental services production. *Health Services Research*, 21(5), 635–661. <https://pmc.ncbi.nlm.nih.gov/articles/PMC1068979/>.
- Sorso. (2026). *Dental Practice EBITDA Multiples 2026: 5x–12x + Calculator*. Reviewed 21 July 2026. <https://www.thesorso.com/answers/dental-practice-ebitda-multiple>. Accessed 28 July 2026.
- Williams, J. B. (1938). *The Theory of Investment Value*. Harvard University Press.

SOURCE-QUALITY NOTE

FOCUS, CT Acquisitions, DDSMATCH, Lakeshore and Sorso are advisor or commercial publications. Their ranges are used as directional market context. Sorso cites FOCUS, so the publications are not treated as independent transaction datasets. No multiple in this paper is represented as an appraisal coefficient. The ADA workbook is retained in this list because it is discussed in Section 7.4; no band in this edition is drawn from it.

Appendix A: Parameter and evidence book

Parameter or rule	Value	Source / basis	Status
Model horizon	5 years	Finite scenario design	Judgment
Contractual window	5 years, horizon-capped	Mechanism assumption	Judgment
Repriced window	3 years	Mechanism assumption	Judgment
Behavioural window	1 year	Mechanism assumption	Judgment
Haircuts	0.90 / 0.70 / 0.40	Evidence and execution rules in Section 4.6	Judgment
Rate scenarios	8%, 12%, 15%; 25% stress	No client hurdle documented	Scenario, not estimate
Switching costs S_i	Not available	Required deduction for net RMV	Case reports gross RMV only
Internal reference	One comparable site	Avoid synthetic component frontier	Specification
Historical duration	6-month average; 24-month history	Rule retained for calibration testing	Judgment
Overhead tolerance	2% of reported overhead	Extraction control	Judgment
External component bands	Not used	Not ADA-sourced; staff band definitionally mismatched (7.4)	Withdrawn (Appendix D)
Fee-schedule line	\$52,000 (median basis)	Restated from \$160,000 75th-percentile basis (5.3)	Restated (Appendix D)
5–8x transaction range	Directional illustration	Advisor publications; subject fit not proven	Conditional sensitivity

Table A1. Every model parameter used in the case is shown with status.

Appendix B: Display-level reconciliation

Quantity	Reconciliation
Candidate + context	$\$185,000 + \$247,000 = \$432,000$
Context restatement	$\$355,000 - \$160,000 + \$52,000 = \$247,000$
Gross RMV at 8%	$\$81,179 + \$35,934 + \$171,377 + \$107,803 + \$27,060 = \$423,353$
Annuity factors at 8%	$A(3) = 2.5771; A(5) = 3.9927$
Contractual haircut run-rate	$(\$10,000 + \$30,000) \times 0.90 = \$36,000$
EVF sensitivity	$\$36,000 \times 5 = \$180,000; \$36,000 \times 8 = \$288,000$
Stock-flow laboratory bridge	\$96,000 stock gap and \$95,000 flow candidate overlap at displayed rounding; not additive
Partial realised margin	$\$4,850,000 - \$800,000 - \$1,984,000 = \$2,066,000$
Partial attainable margin	$\$5,160,000 \times 0.97 - \$704,000 - \$1,856,000 = \$2,445,200$
Omitted current cost proxy	$\$2,066,000 - \$920,000 = \$1,146,000$
Internal MCR upper bound	$\$920,000 / (\$2,445,000 - \$1,146,000) = \text{approximately } 70.8\%$
Fee restatement	$\sim \$73,000 \text{ median-basis production} \times 0.712 \text{ PPO-collectible ratio} = \sim \$52,000$
Buyer-style scenario at $h = 0.50$	$\$45,099 + \$19,964 + \$122,412 + \$59,891 + \$19,328 = \$266,694; \$266,694 / \$423,353 = 0.630$

Table B1. This reconciles displayed arithmetic, not confidential source evidence.

Appendix C: Model governance and disclosure

Model risk	Control	Analytical effect
Incomplete cost denominator	Require a complete partition or restore omitted costs under an explicit assumption	MCR is reported as an upper bound until a full P&L is available
Incompatible frontier inputs	Separate the internal reference-site measure from external stress cases, and check that any external band measures the same quantity as the extracted line	Primary and stress results remain distinct; mismatched bands are withdrawn rather than caveated
Ambiguous formula bases	Define B_c as the documented component-specific base	Formula units can be audited
Benchmark target drift	Record which percentile or reference each opportunity is measured against, in the same table as the amount	A median claim cannot inherit a 75th-percentile figure
Unestimated hurdle rate	Report multiple discount-rate scenarios without a conservatism label	Rate dependence remains visible
Unproven persistence	Condition EVF on sustained, accepted normalised EBITDA	Transaction sensitivity is not presented as permanent value
Unquantified switching costs	Deduct S_i when observed and label the case scenarios gross	Net RMV is not claimed
Counterparty discretion	Apply owner-action and counterparty-dependency gates, and report gate status counts	Register candidacy remains conditional and its conditionality is quantified
Opportunity described as realised cash	Separate exposure, modelled opportunity and one-time recovery	Counterfactual amounts are not called cash loss
Unsupported external component bands	Withdraw bands that are unsourced or definitionally mismatched	The paper carries no external comparator
Register described as a pipeline	Report the register by execution route and count the resolved action gates	The owner-only tier is visible and separate from the one-agreement-away tier
Confidential case evidence	Disclose the data boundary and propose a synthetic replication package	Readers can distinguish author-reported from reproducible evidence
No implementation follow-up	Report outcomes as not observed and make no realised-recovery claim	Retrospective status is not treated as validation
Single-author verification	Obtain independent review of a sample of classifications and calculations before publication	Two errors in the first edition survived internal review (7.10)

CONFLICT AND FUNDING DISCLOSURE

The author is the founder of Marginr, which sells the margin audit described in this paper. The engagement analysed is a Marginr client, used with the owner's consent and anonymised. The retrospective analysis was provided to the owner at no charge. No third party funded this work. The author designed the model, performed the engagement and conducted the retrospective analysis.

DATA AVAILABILITY

Client source documents are confidential and are not available with this paper. A future reproducibility package should include a synthetic case, data dictionary, parameter file, calculation code and test outputs. Until then, the case evidence should be read as author-reported.

PROFESSIONAL-USE NOTICE

DMF is an operating diagnostic and scenario framework. It is not an appraisal, fairness opinion, investment recommendation, accounting opinion, legal opinion or guarantee of recovery or transaction value.

Appendix D: Revision record

Edition	Date	Change	Effect on results
Ed. 1	July 2026	First publication.	—
Ed. 2	August 2026	The fee-schedule line was restated from the 75th percentile to the median (\$160,000 to \$52,000). The identified total moved from \$540,000 to \$432,000 and the context total from \$355,000 to \$247,000. The labour line was relabelled as dentist-inclusive. The external band comparison and its stress case were withdrawn.	Register, RMV, EVF and internal MCR unchanged. Both corrections sat in lines the partition had already excluded from claim-bearing outputs.
Ed. 3	August 2026	Executive summary added. It presents the case in tiers: Tier 1, the recoverable register, with an owner-only sub-tier; and Tier 2, conditional opportunity. Figure 1 presents the identified total as a composition rather than a descending sequence. Register execution-route tiers made explicit (5.4.1). Switching-cost breakevens (5.6.1) and a buyer-style realisation scenario (6.3) added. Stock-to-flow bridge added (5.7.1). Validation programme restructured with instruments and timing (Section 8), and its record schema published as Appendix E. Rights, citation and permitted-use page added. Identifying details further reduced.	No numerical result changed.

Table D1. Every change between editions, with its effect on the paper's results.

Internal re-reading found both second-edition errors, and no external check produced them (Section 7.10). The validation programme in Section 8 therefore treats independent review of source classifications as a requirement, and Appendix C carries single-author review as a named model risk.

Appendix E: Follow-up record schema

This schema is the preregistered data dictionary for the validation programme in Section 8. One record is written per finding, at day 30 and again at day 90, and then at each persistence measurement. The schema is published before data collection begins, so that the fields cannot be selected after outcomes are known.

Field	Type	Definition	Rule
Finding id	Identifier	Stable key linking the record to one register or context item	Assigned at the reporting date; never reused
Implemented	Yes / no / partial	Whether the action was taken by the measurement date	Partial requires the implemented share to be stated
Date	Date	Date the action was taken, or the date of the measurement if not taken	Recorded from dated evidence, not from recall
Realised annual \$	Currency per year	Recovery measured on comparable volume, pre and post	Zero is a valid value and is recorded as zero
Evidence reference	Document reference	Invoice, statement, contract or receipt that supports the amount	A record without a reference counts as unevidenced
Reason if abandoned	Coded reason	Why the action was not taken, or was reversed	Abandoned findings are retained, never deleted
Censoring status	Active / failed / complete	Whether the item was still running at the measurement date	Active items are right-censored, not assumed persistent or failed

Table E1. The record schema for Section 8. It applies to every finding, including those that were never implemented.

Records are held per engagement and reported in aggregate. Development cases and validation cases are held in separate sets. An independent reviewer checks the source classifications, sampling at least 20% of register items, before any outcome is published.

Rights, citation and permitted use

Copyright

The text, tables, figures and layout of this paper are the author's original work. They are protected by copyright from creation, without registration, under the Berne Convention and the French Code de la propriété intellectuelle. Copyright protects this expression. It does not protect the underlying method, which is published deliberately.

Licence

This paper is published under the Creative Commons Attribution-NoDerivatives 4.0 International licence (CC BY-ND 4.0), at <https://creativecommons.org/licenses/by-nd/4.0/>. You **may** copy and redistribute this paper in any medium or format, including in a commercial context, provided you credit the author, link the licence, and distribute it **whole and unchanged**. You **may not** distribute a modified, adapted, abridged or rebranded version, and you may not present the method under another name as your own work. The licence grants no rights in the trade marks.

Trade marks

Marginr™, Marginr Research™, Discounted Margin Flow™ and DMF™ are unregistered trade marks of Paul-Matthieu Castro. No trade mark rights are licensed by this paper. Redistribution of the paper conveys no right to use these names for any service or publication.

Reserved and confidential material

The method is published. The implementation is not. The source code, the calibrated parameter values, the extraction procedures, the benchmark source records and all client evidence files are confidential. This paper does not licence or disclose them. They are maintained as trade secrets under Directive (EU) 2016/943 and the French law transposing it, with documented access restrictions. The synthetic replication package committed to in Section 8 is the route by which a third party may check the mechanics without access to confidential material.

How to cite

Castro, P.-M. (2026). *Discounted Margin Flow: developing a new transparent model of margin health for dental practices, with a retrospective application to a three-site group*. Marginr Research Paper 2026-07, third edition. Lyon, France. <https://doi.org/10.5281/zenodo.21907533>

First publication record: first edition 28 July 2026; second edition 10 August 2026; third edition 12 August 2026. This edition is archived under the digital object identifier above, which resolves to the deposited record.

No warranty and no reliance

This paper is published for information. It is not financial, legal, tax, accounting or investment advice. It is not an appraisal, and it is not a quality-of-earnings review. The figures relate to one anonymised engagement and are not representative of any other practice. No reader should act on this paper without independent verification. The author accepts no liability for any decision taken in reliance on it. Appendix C states the conflict, funding and data disclosures.

Third edition, dated 12 August 2026. It supersedes the editions dated 28 July 2026 and 10 August 2026. URLs and access dates are included for web sources. All displayed values remain rounded, and percentages derived from rounded values are approximate.