MYLEVEL™: A PHARMACOKINETIC FRAMEWORK AND CLINICIAN TOOL FOR PRECISION GLP-1 DOSING

A Clinical Whitepaper

Ian Justl Ellis, MD

Founder & CEO, MyLevel Inc. | Voafit

I. EXECUTIVE SUMMARY

GLP-1 receptor agonists represent one of the most significant pharmacologic advances in recent history. As these medications have moved into widespread real-world use, however, serious problems have emerged. Real-world discontinuation rates approach 50–65% within the first year,[1–2] driven primarily by side effects, dose intolerance, and cost.[3] Among patients who remain on therapy, many experience clinically meaningful negative consequences: disproportionate lean mass loss,[4–5] reductions in bone mineral density,[6–7] chronic gastrointestinal symptoms,[8–9] and nutritional compromise.[8]

Many of these problems represent not an intrinsic failure of the medications, but a predictable outcome of how they are dosed in current protocols.

The pharmacokinetic basis of the dosing problem is straightforward. GLP-1 medications have half-lives of approximately five to seven days,[10–11] causing prior doses to additively contribute to total exposure for weeks until steady state is reached in approximately 4–5 weeks.[10–12]

A patient's day-to-day experience on a GLP-1 therefore depends primarily on cumulative medication level — not the milligrams injected in a given week.

Restated plainly: It's not the dose that matters most. It's the level.

The landmark GLP-1 trials used elegant escalation protocols to progressively increase medication levels toward predetermined maximum doses.[13] This proved population-level efficacy, but when applied to individuals in routine clinical practice, the result is relentless escalation toward a fixed ceiling with no mechanism to identify or maintain an individual therapeutic window.[13–14]

The MyLevel™ dosing engine was created to solve this problem. MyLevel uses published pharmacokinetic constants and actual dose history to estimate the patient's real-time medication level and calculate the precise dose needed to stay within their identified therapeutic range. Providers adjust only two inputs — target peak medication level and dosing interval — which together define the patient's exposure window.

Observational data from 112 patients across approximately 2,800 patient-weeks treated under this framework demonstrate weight loss rates meeting or exceeding Phase III trial predictions at approximately half the cumulative medication exposure, with zero treatment discontinuations attributable to side effects. These comparisons are indirect, benchmarked against published trial trajectories rather than a concurrent control arm, and are presented as hypothesis-generating.

II. THE PHARMACOLOGIC BASIS AND FAILURE OF STANDARD GLP-1 DOSING

After administration, each dose of a GLP-1 receptor agonist produces a concentration curve that rises following absorption, peaks, and then declines gradually as the drug is eliminated (FIGURE 1).[10–11]

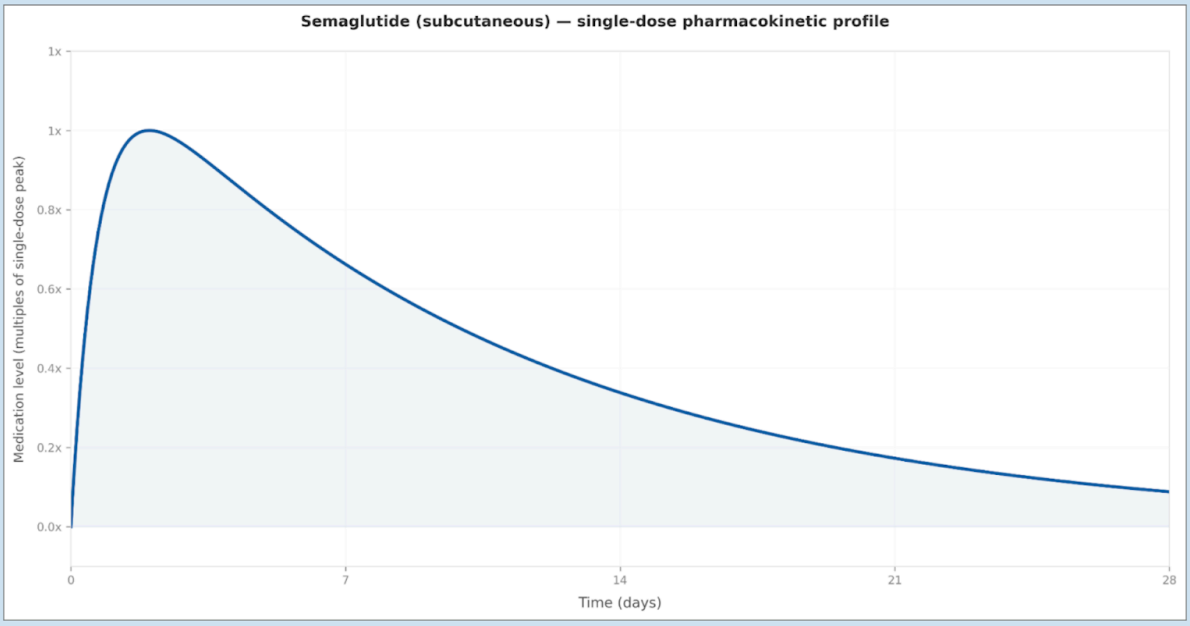

FIGURE 1. Single-dose pharmacokinetic profile of subcutaneous semaglutide. *Following injection, medication level peaks around day 2 and declines with a half-life of approximately 7 days.*

Because elimination occurs slowly relative to the dosing interval — with half-lives of approximately five to seven days — each new dose is added to the remaining contribution of prior doses that have not yet cleared.[12] Over time this produces a composite exposure curve reflecting the combined pharmacologic contribution of multiple recent injections. With repeated dosing, medication levels rise after each administration, reach a peak, and then decline until the next dose. The peak and trough together define the patient's exposure range across the dosing interval (FIGURE 2).[10–12]

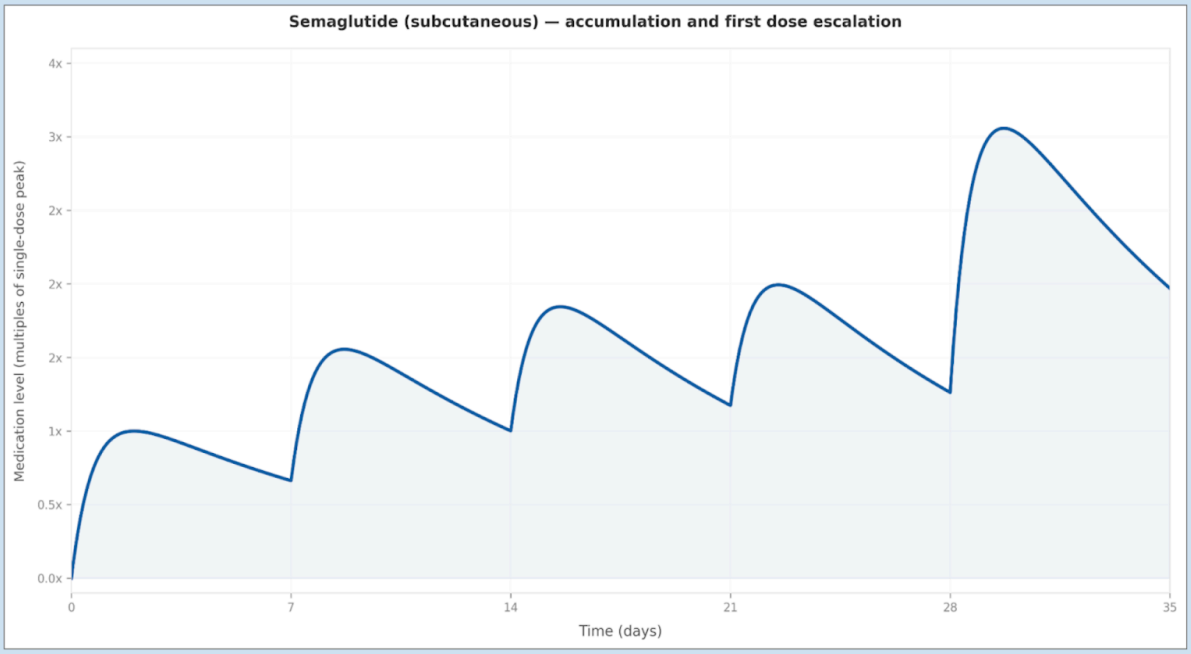

FIGURE 2. Accumulation and dose escalation. *Four identical 0.25 mg weekly doses of semaglutide produce progressively higher peaks as prior doses have not fully cleared. The scheduled doubling to 0.5 mg on dose 5 produces an immediate jump to approximately 3× the initial peak. Five weeks of treatment, five different medication levels.*

These dynamics explain several common clinical observations that are frequently misattributed to other causes. Patients report the same dose feels stronger after several weeks.[10,12] Side effects appear weeks into treatment on a stable dose.[8] Hunger control varies across the dosing interval. Missed or delayed doses alter the experience

unpredictably — because the underlying medication level has shifted in ways neither patient nor clinician can see.[11]

Medication level is among the variables most directly linked to a patient's clinical experience — and the one least visible under current prescribing practice. Dose functions primarily as the input used to shape that exposure pattern.

The Therapeutic Window Problem

The standard escalation schedule was developed to move large trial populations toward predefined maintenance doses to show strong average outcomes at the group level.[13–14]

This was useful in a trial setting, but it was not designed to ask the most clinically relevant question:

What medication level range is actually needed for an individual patient to achieve effective appetite control with acceptable tolerability?

A therapeutic window is bounded below by the minimum effective drug level and above by the level at which adverse effects become unacceptably likely.[16] Clinical experience strongly suggests that for GLP-1 medications, the breadth and location of this window vary substantially between individuals.[17] Treatment within the ideal window means consistent appetite control with manageable side effects. Treatment below it means insufficient appetite suppression. Treatment above it means escalating risk without additional benefit (FIGURE 3).

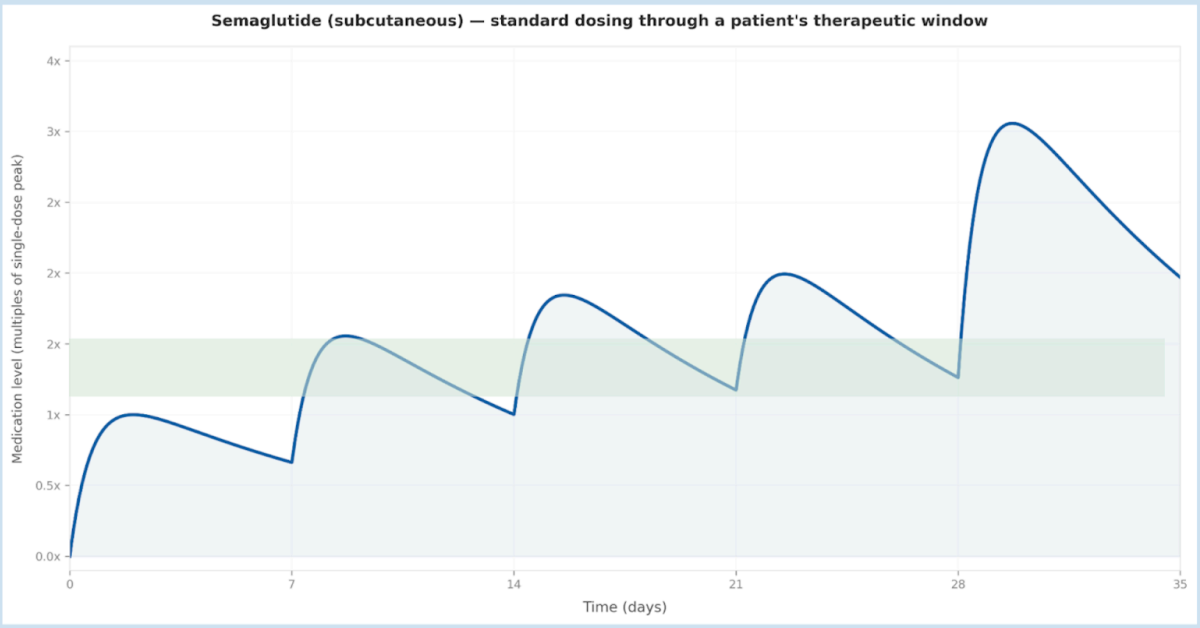

FIGURE 3. Standard semaglutide dosing through a patient's therapeutic window. *The exposure curve generated by dose #1 is subtherapeutic. The peak reached by dose #2 stays within the patient's therapeutic window. However, the next dose is taken after the subsequent trough dips too low. Subsequent doses result in peaks exceeding the patient's tolerance.*

The Consequences of Progressive Escalation

There is a fact about standard GLP-1 prescribing that is rarely stated directly: every clinician who follows the standard dosing protocol is progressively, continuously increasing their patient's GLP-1 exposure with each successive dose.[10,12] This escalation is a feature and legacy of the trial design (FIGURE 4).[13–14]

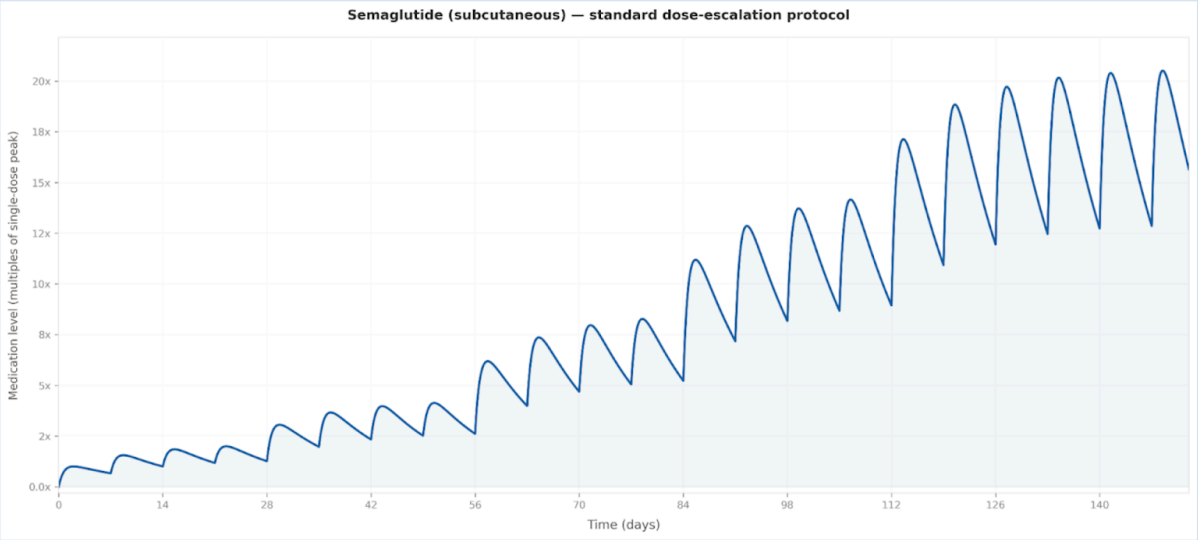

FIGURE 4. Standard semaglutide dose-escalation protocol. *Medication levels rise continuously across five dose phases (0.25, 0.5, 1.0, 1.7, 2.4 mg weekly) lasting four weeks each, reaching approximately 20× the initial single-dose peak at steady state.*

The exposure increase is obvious at the time of explicit dose increases. But the standard protocol's recurring pattern of four stable doses followed by a scheduled increase means that patients in between those increases *also* experience progressively higher medication levels as accumulation develops.[10,12] By the time steady state arrives, the next escalation occurs, and the process begins anew. The climb to the maximum dose is continuous.

A patient following the study protocol perfectly would not have a consistent, stable week-to-week medication experience until after steady state is reached at week 20 — after which point there is nowhere further up to go.

This might work well for some individual patients. But for many, it is a recipe for disaster.

Because the standard dosing ladder has no reliable method for identifying an individual patient's therapeutic window, patients whose exposure rises above their tolerable range

either continue escalating and suffer through unnecessary side effects, or discontinue treatment altogether.[1–3] For patients who push through, sustained excessive exposure produces the constellation of adverse outcomes described in the preceding summary — not as rare idiosyncratic reactions, but as expected consequences of sustained pharmacologic overexposure.[4–9]

Additionally, every unnecessary escalation step consumes therapeutic headroom that cannot be easily recovered. A patient who arrives at maximum dose without ever identifying their minimum effective level has no good option when progress eventually slows.

Clinicians can hold a patient at a fixed dose longer than 4 weeks, but stabilization is available only at a small number of predefined dose levels determined by trial design, not individual response.[13–14] A patient who finds their effective level between two standard doses has no way to stay there. Some providers escalate in smaller increments, but fixed dosing of any size still results in progressive accumulation and an inconsistent experience.[10,12]

A more precise framework would identify each patient's therapeutic window, maintain treatment within it, and adjust only when clinically indicated (FIGURE 5).

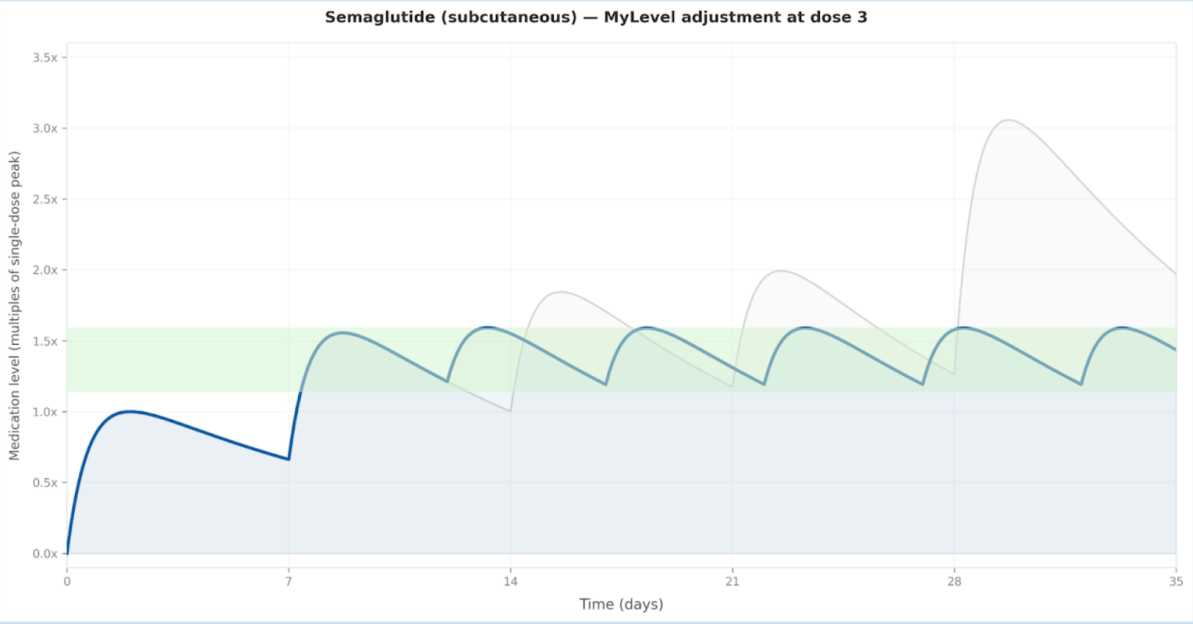

FIGURE 5. Idealized protocol adjusting to a patient's therapeutic window. *Dose #1 is again subtherapeutic. Dose #2 resulted in excellent appetite suppression without side effects, but the patient's hunger returned before the subsequent scheduled dose. Real-time adjustments beginning with Dose #3 (taken when the patient's hunger returned) result in peaks and troughs staying within the therapeutic window.*

III. THE MYLEVEL FRAMEWORK

The most effective and truly personalized approach to GLP-1 therapy would need to answer five clinical questions that standard dose-escalation protocols do not address.

1. What is the ideal target peak medication level for this patient?
2. What dose amount is required at the moment of administration to consistently reach that target?
3. What is the trough level below which the patient no longer experiences beneficial effects?
4. How frequently should medication be taken to ensure levels do not fall below the effective range?
5. How should target level and dosing interval be adjusted over time?

MyLevel is both a framework and a patent-pending technological toolset designed to answer these questions.

How It Works

The framework turns on a simple inversion in treatment logic. Rather than prescribing a fixed weekly dose that results in escalating medication levels,[10,12] the clinician selects a target peak medication level (MyLevel) and desired dosing interval. Together, these two variables fully define the patient's exposure window.

The dosing engine uses published pharmacokinetic constants[10–12] to continuously calculate the patient's real-time estimated medication level, determine the distance to the target, and compute the precise dose required at any moment to reach it. The uncontrolled accumulation inherent in fixed-dose escalation protocols is eliminated from the outset.

For ease of communication, medication levels are mapped onto a normalized 0–50 scale. A MyLevel of 0/50 represents the absence of active medication; 50/50 corresponds to the steady-state level produced by the highest standard dosing regimens in major trials.[13–14] Each integer represents a 2% step — delivering approximately 10-fold greater granularity than the standard dosing ladder. For reference, a dose increase from 2.5 mg to 5 mg weekly tirzepatide is equivalent to 9 MyLevel step increases.[14]

Adjustment is intuitive: raise the target for more appetite suppression, shorten the interval to raise the trough without changing the peak, or combine both (FIGURE 6). Adjustments can also be made in the opposite direction.

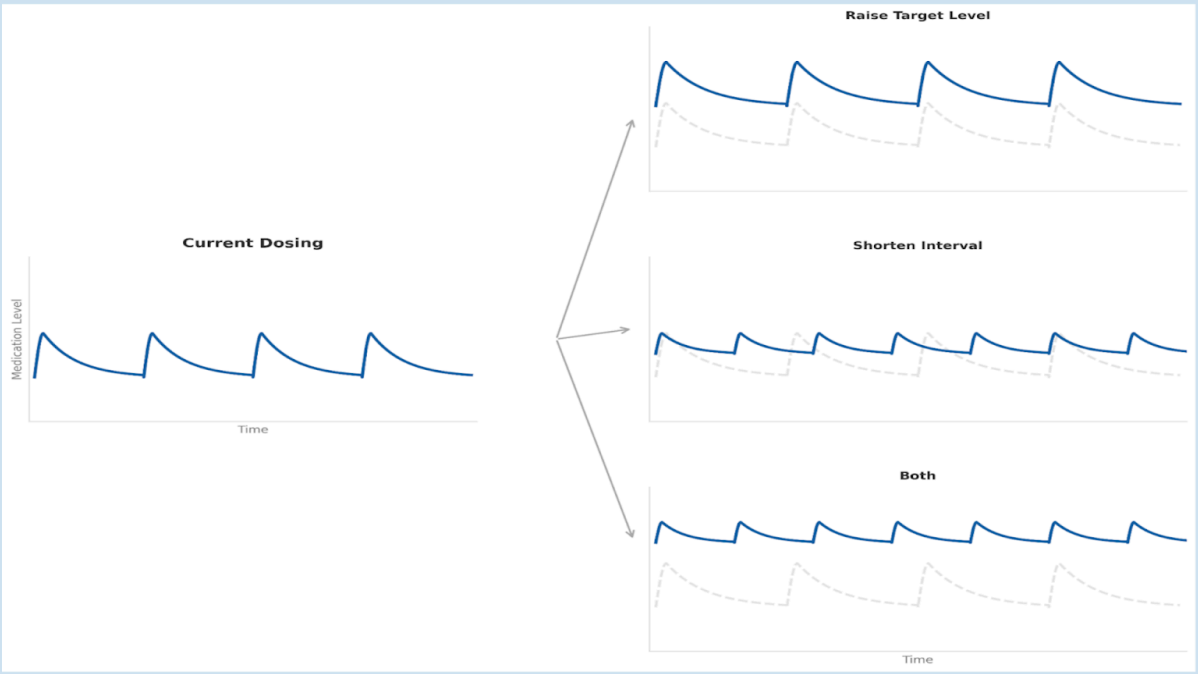

FIGURE 6. The effects of the two controls available within the MyLevel framework. *From a baseline dosing pattern, three adjustments are shown: raising the target level shifts the entire oscillation envelope upward without changing its shape; shortening the dosing interval maintains the same peak while raising the trough, compressing the oscillation amplitude; combining both controls shifts the envelope upward and compresses it simultaneously.*

MyLevel converts GLP-1 dosing from a fixed escalation ladder into a controllable, individualized exposure-management system — one designed to find and keep each patient within their own therapeutic window[16] throughout the entire course of therapy.

Modeling Limitations and Sources of Variability

The medication levels generated by MyLevel are pharmacokinetic estimates, not direct serum measurements. The model relies on published population-derived constants[10–12] and assumes consistent adherence and accurate dose logging. Individual patients may vary due to differences in metabolism, body composition, injection site, and absorption characteristics.[17] This approach parallels established practice in therapeutic drug monitoring, where modeled exposure estimates — such as vancomycin AUC calculations or cyclosporine trough-based dosing — guide clinical decisions without requiring

continuous direct serum measurement. The value of modeled level lies not in its absolute precision but in its ability to make relative exposure visible and adjustable in real time.

IV. CLINICAL APPLICATION

Initiation and Adjustment

Treatment typically begins at a target of MyLevel 1/50 — approximately two percent of the total exposure range — allowing early tolerability assessment. Patients with known GLP-1 tolerance may start higher. The dosing engine calculates a loading dose to reach the target immediately, producing a stable exposure pattern from the outset rather than a gradually rising accumulation curve.[10,12]

Subsequent adjustments follow a simple test: if appetite suppression is insufficient, increase the target; if appetite control and tolerability are ideal, hold; if exposure feels excessive, reduce. Weekly asynchronous check-ins capture hunger level, protein intake, side effects, exercise activity, and the patient's subjective assessment of whether their current level feels too low, appropriate, or too high. These responses are displayed longitudinally on a provider dashboard.

Importantly, a stalled patient does not always need a medication adjustment. A patient whose appetite is only a 1 on a 5-point scale but who has not been exercising does not need more medication — they need to move more. Without this longitudinal context, a stall might prompt unnecessary escalation when the real issue lies elsewhere entirely.

When increases are indicated, a single MyLevel integer represents a 2% step — far smaller than the 50–100% dose increases typical of early standard escalation.[13–14] Increases greater than two units at a time are generally discouraged.[16]

Dosing Interval Adjustment

Adjusting the target level addresses only the upper boundary of exposure. The lower boundary is determined by the time between doses. Patients with narrow therapeutic windows may experience both side effects near peak and breakthrough hunger near trough within the same cycle. Standard protocols address low troughs by increasing the dose — which simultaneously raises peak exposure past the tolerable ceiling.[13–14] MyLevel separates these variables. A patient whose peak is tolerable but whose trough drops below the effective range can shorten the dosing interval, maintaining the same peak while raising the trough (FIGURE 7).

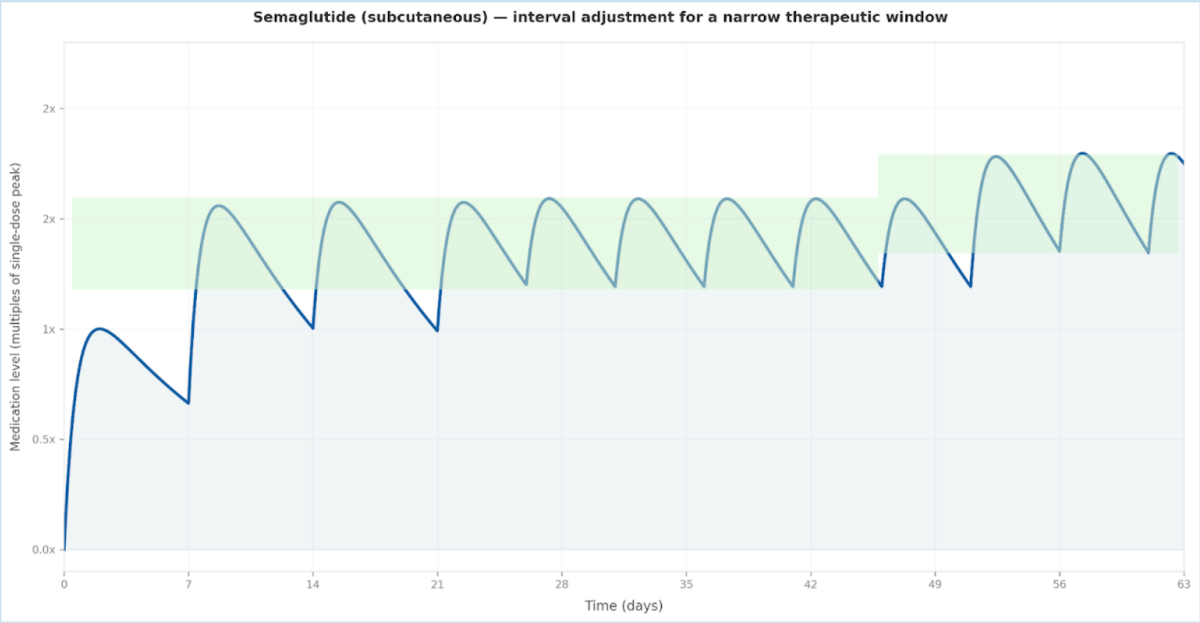

FIGURE 7. Dosing interval adjustment for a patient with a narrow therapeutic window. *The first two doses match the standard protocol. The second dose identifies the ideal peak medication level. However, on a 7-day interval, the trough drops well below the therapeutic floor between doses. Shortening the interval to 5 days maintains the same peak level while raising the trough to remain within the window. A subsequent target level increase demonstrates the system's ability to adjust to changing therapeutic windows over time.*

Optional Doses

The engine also calculates an always-available optional dose allowing the patient to reach their target at any time, independent of the base schedule. Late doses, early doses, temporary increases for high-stress periods, or deliberate brief reductions for dietary flexibility — in each case the system recalculates from the actual current estimated level, preventing overshoot. All optional doses are logged and visible to the supervising provider.

The MyLevel Interface

The provider and patient interfaces are shown in FIGURES 8–11.

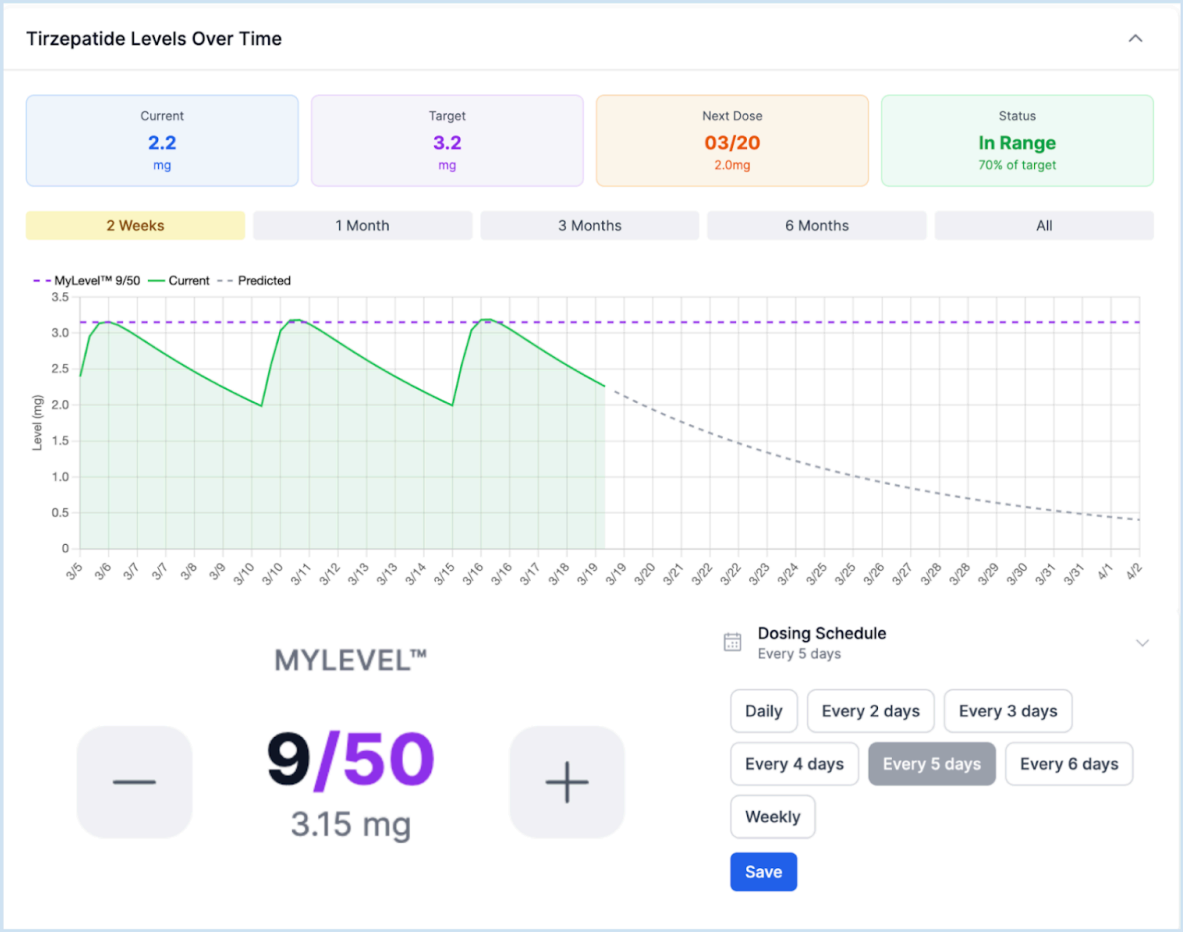

FIGURE 8. MyLevel adjustment provider interface. *The provider sets the target MyLevel (in this case 9/50, for a peak tirzepatide level of 3.15mg) and preferred dosing interval using a simple interface. All patient-facing dosing calculations are calculated automatically by the dosing engine.*

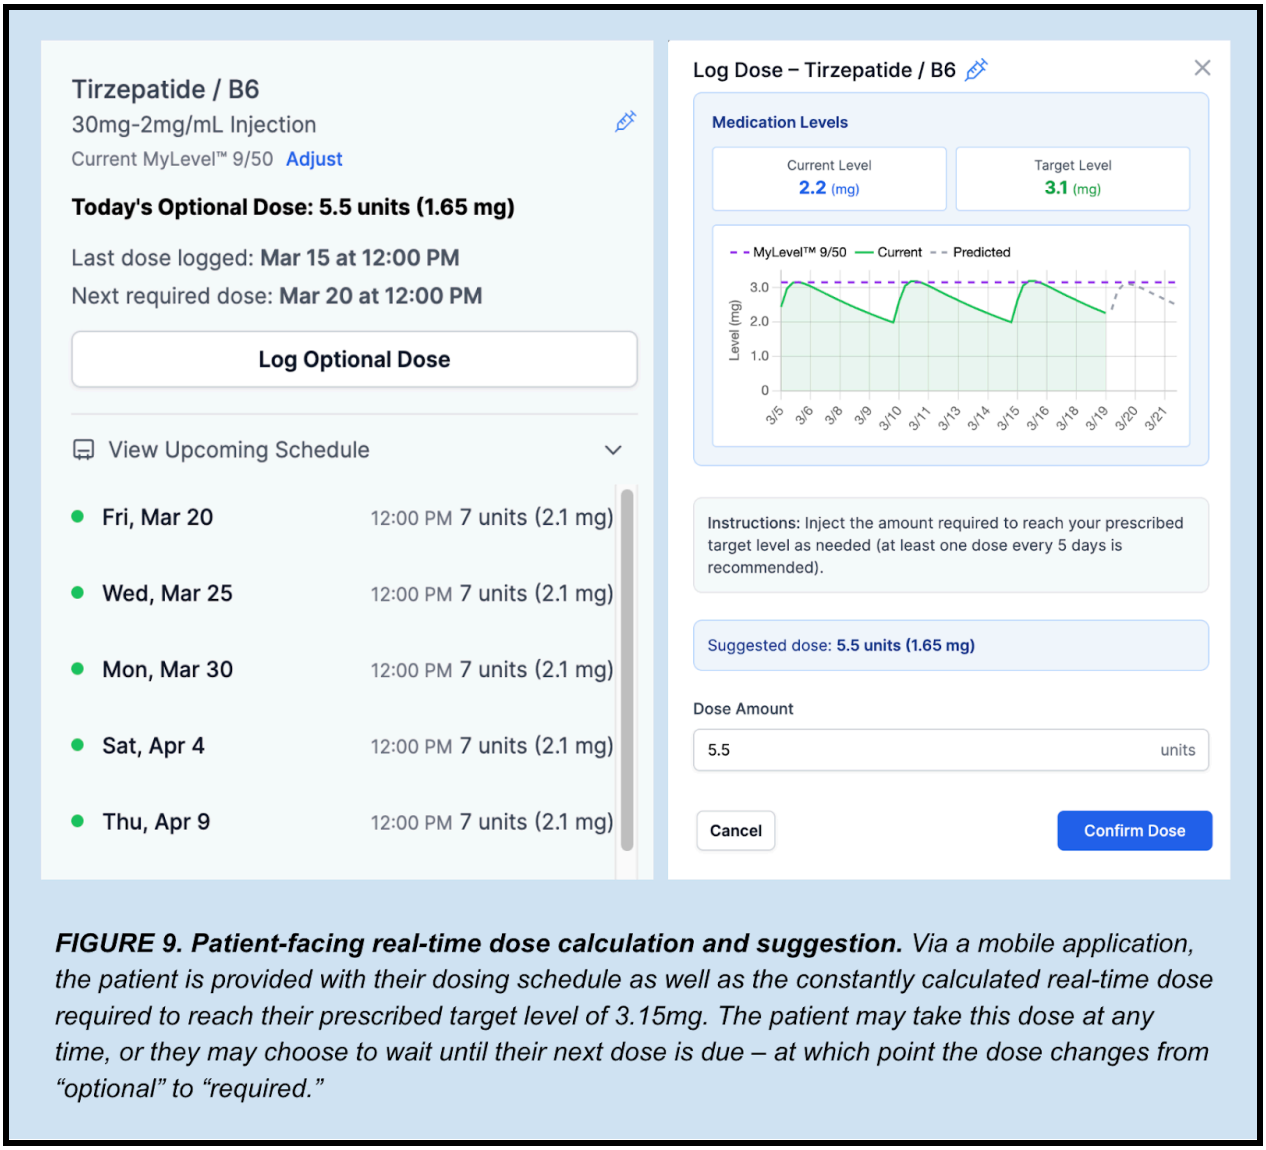

FIGURE 9. Patient-facing real-time dose calculation and suggestion. *Via a mobile application, the patient is provided with their dosing schedule as well as the constantly calculated real-time dose required to reach their prescribed target level of 3.15mg. The patient may take this dose at any time, or they may choose to wait until their next dose is due – at which point the dose changes from "optional" to "required."*

FIGURE 10. Adjustment of MyLevel and dosing interval. *The provider has changed the MyLevel target to 11/50 from 9/50 and adjusted the dosing interval from every 5 days to weekly. The effects of this change are immediately visible on both patient and provider dashboards.*

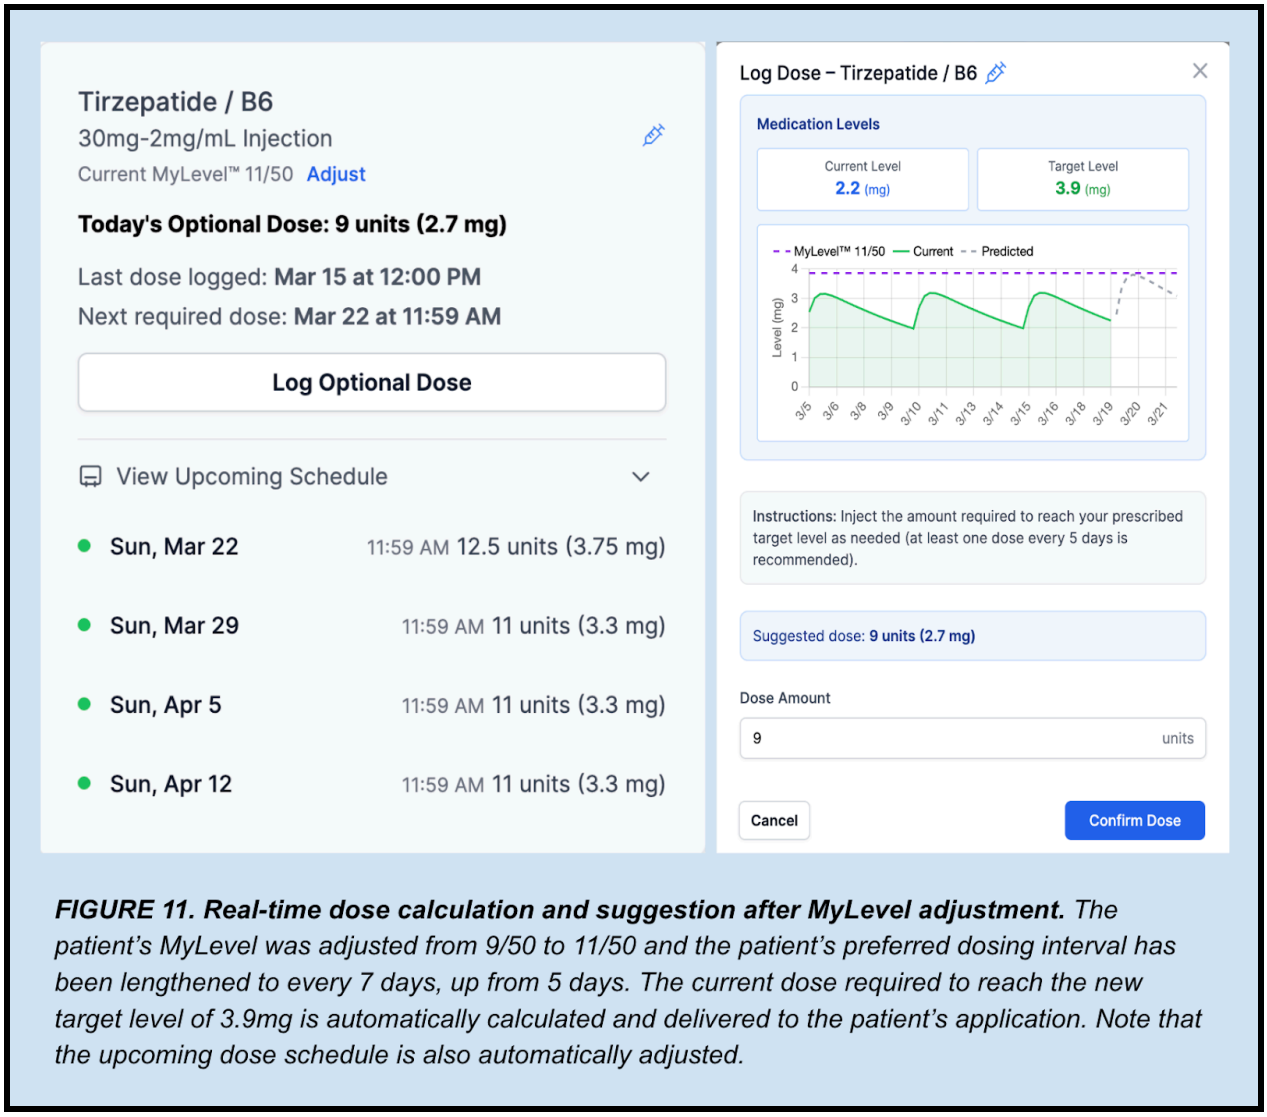

FIGURE 11. Real-time dose calculation and suggestion after MyLevel adjustment. *The patient's MyLevel was adjusted from 9/50 to 11/50 and the patient's preferred dosing interval has been lengthened to every 7 days, up from 5 days. The current dose required to reach the new target level of 3.9mg is automatically calculated and delivered to the patient's application. Note that the upcoming dose schedule is also automatically adjusted.*

The 30-Week View

Viewed over a full 30-week treatment course, the cumulative effect of exposure-guided dosing becomes clear. Under the standard Wegovy escalation protocol, a patient progresses through five dose phases to a steady-state level approximately 20 times the initial single-dose peak by week 20, where it remains. A representative MyLevel patient — beginning at one-third of the standard starting dose and adjusting gradually based on clinical response — reaches approximately half the cumulative medication exposure while medication levels continue to rise through week 30 (FIGURE 12).

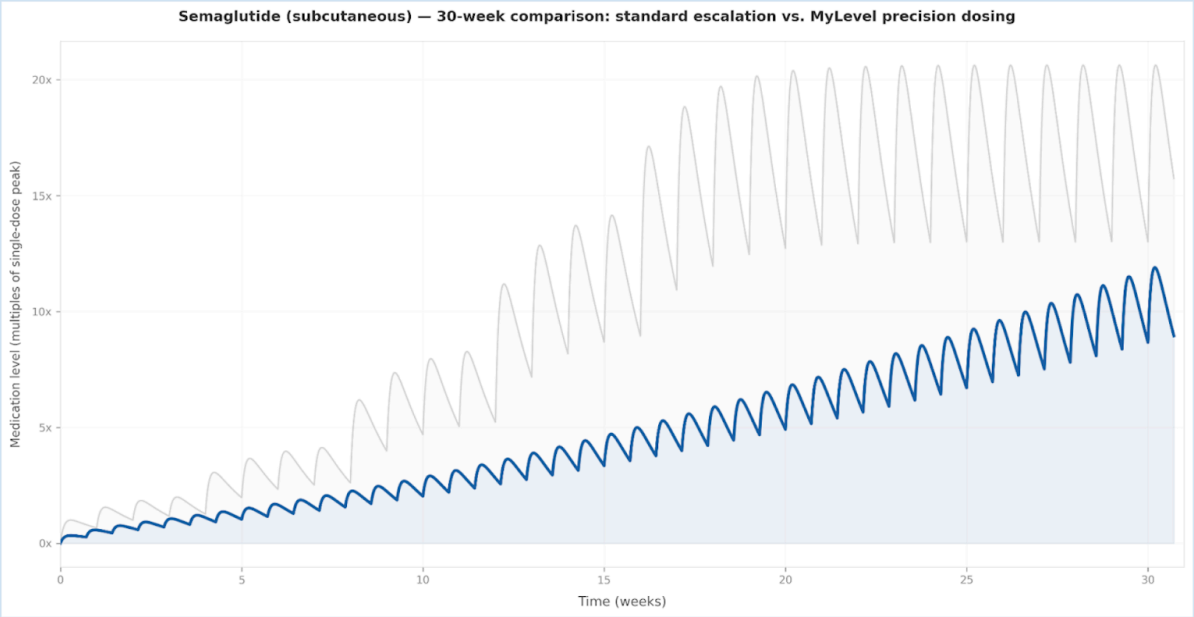

FIGURE 12. 30-week comparison of standard semaglutide escalation (gray) versus MyLevel precision dosing (blue). *The gray curve represents the standard Wegovy escalation protocol through steady state. The blue curve is a representative exposure trajectory modeled from the semaglutide cohort's observed average weekly dose of 0.71 mg, beginning at one-third of the standard starting dose and titrating gradually based on the cohort's mean escalation pattern. This is an illustrative model, not a direct trace of any individual patient's exposure. Standard dosing reaches its ceiling at week 20; the MyLevel trajectory continues to rise through week 30 at approximately half the cumulative medication consumption.*

V. ILLUSTRATIVE CASES

The following cases demonstrate the range of presentations the MyLevel framework is designed to accommodate — and why aggregate outcomes cannot be understood without appreciating the individual variability that produced them.

Case 1 — The Author. The author's own initial semaglutide course on the standard escalation schedule[13] produced 25 pounds of weight loss over three and a half months, but body composition analysis revealed that 16 of those 25 pounds were lean mass — substantially worse than achieved through diet and exercise alone.[4–5] Appetite suppression had progressed to near-complete food aversion by week three on the same dose, with persistent GI symptoms and declining training capacity. A subsequent course managed under the exposure-guided framework achieved comparable weight loss with

markedly better body composition (fat loss significantly exceeding lean mass loss) and absent side effects.

Case 2 — Medication-Sensitive Patient. A young patient seeking to lose 25 pounds reported significant prior medication sensitivity. Tirzepatide was initiated at one-third of the standard starting dose.[14] Despite this reduction, severe GI symptoms occurred following the first injection.[8] Under the standard protocol, the options would be to wait (while accumulation continued to drive levels upward[10–12]) or discontinue. Under MyLevel, the estimated level was calculated and the target was reduced before the next dose. Side effects did not recur. The patient completed treatment successfully, reached her goal weight, and never required a dose exceeding the standard starting dose.

Case 3 — High-Complexity Patient. A patient in her mid-sixties (BMI ~60, severe osteoarthritis, cardiovascular disease, recent sepsis with acute kidney injury) could not safely tolerate even one week of significant adverse effects. Tirzepatide was initiated at a fraction of the standard starting dose with deliberately slow titration over several months. At sixteen months, she had lost 92 pounds, successfully underwent knee replacement surgery, and experienced no medication-related adverse events throughout. The standard ladder does not have a mechanism for the degree of caution her case required.

Each patient required a meaningfully different minimum effective level. Each outcome depended on the ability to identify and maintain that level rather than escalating toward a predetermined dose.[17]

VI. OBSERVATIONAL OUTCOMES FROM THE VOAFIT COHORT

Outcomes were analyzed from 112 patients treated at Voafit under the MyLevel dosing framework since April 2024, representing approximately 2,800 patient-weeks (TABLE 1). The cohort ranged in age from 16 to 79 (mean 42), with a mean BMI of 31.7. Forty-six percent had a BMI below 30 at initiation — below the thresholds typically required for trial enrollment.[13–14]

Because of the lower starting BMIs, the primary efficacy metric is weight loss rate (pounds per week) compared time-matched and weight-matched against Phase III trial predictions.

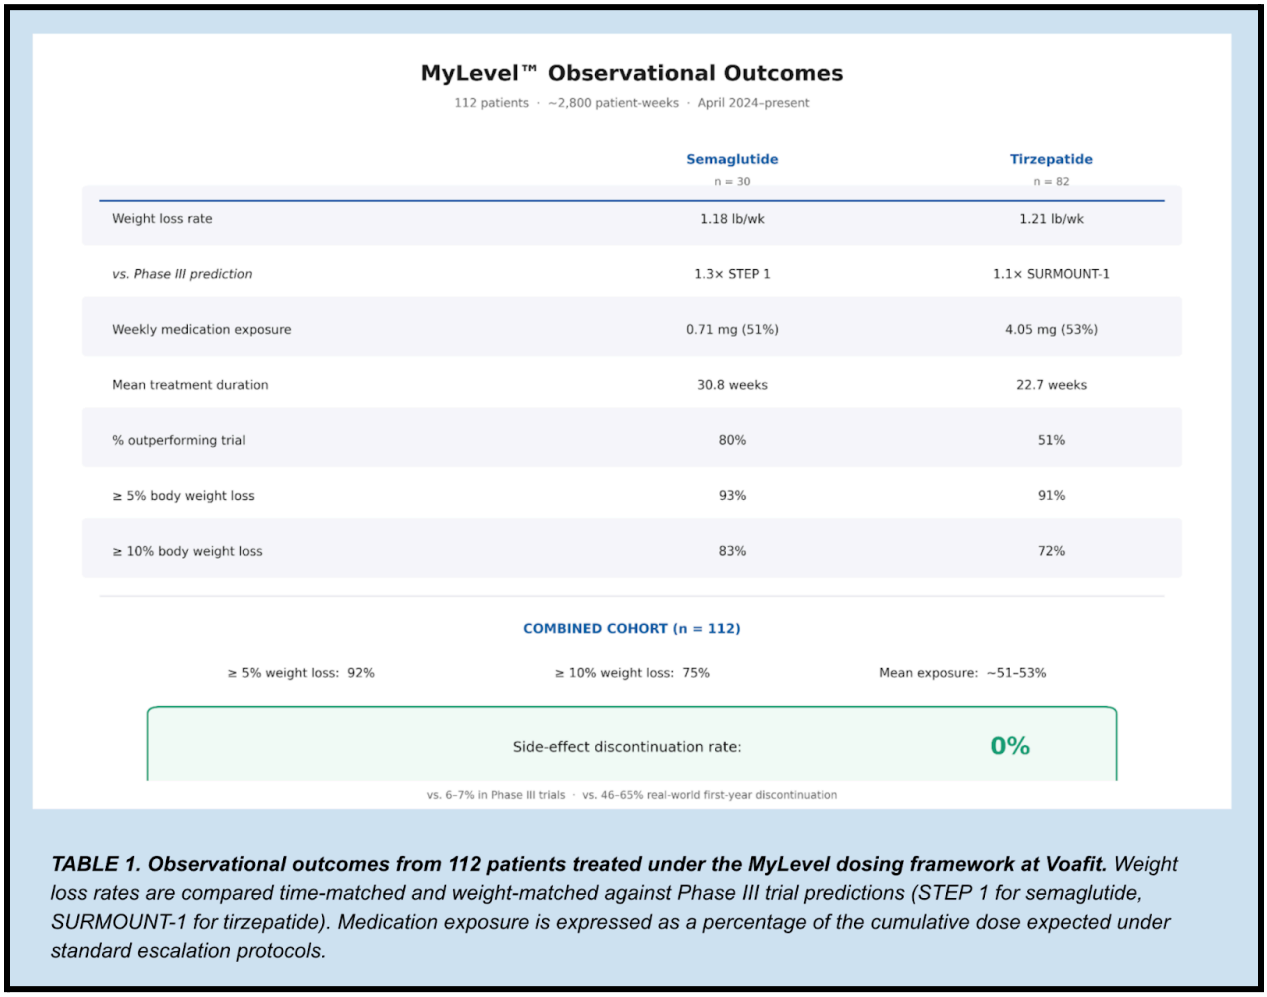

MyLevel™ Observational Outcomes

112 patients · ~2,800 patient-weeks · April 2024–present

	Semaglutide n = 30	Tirzepatide n = 82
Weight loss rate	1.18 lb/wk	1.21 lb/wk
vs. Phase III prediction	1.3× STEP 1	1.1× SURMOUNT-1
Weekly medication exposure	0.71 mg (51%)	4.05 mg (53%)
Mean treatment duration	30.8 weeks	22.7 weeks
% outperforming trial	80%	51%
≥ 5% body weight loss	93%	91%
≥ 10% body weight loss	83%	72%

COMBINED COHORT (n = 112)

≥ 5% weight loss: 92% | ≥ 10% weight loss: 75% | Mean exposure: ~51–53%

Side-effect discontinuation rate: **0%**

vs. 6–7% in Phase III trials · vs. 46–65% real-world first-year discontinuation

TABLE 1. Observational outcomes from 112 patients treated under the MyLevel dosing framework at Voafit. *Weight loss rates are compared time-matched and weight-matched against Phase III trial predictions (STEP 1 for semaglutide, SURMOUNT-1 for tirzepatide). Medication exposure is expressed as a percentage of the cumulative dose expected under standard escalation protocols.*

Cohort and Methods

Patients treated at Voafit under the MyLevel dosing framework between April 2024 and the data cutoff were included if they met three criteria: consistent dose logging within the MyLevel system, maintained weekly weigh-ins, and no obvious confounding therapies during the active weight loss period. One hundred and twelve patients met these criteria. The analysis covers the period from each patient's initiation of

exposure-guided dosing through their lowest recorded weight, representing the active weight loss phase of treatment. Weights were obtained via patient-reported home scale measurements recorded at each weekly check-in.

Each patient's weight loss rate (pounds per week) was compared against the trajectory predicted by Phase III trial data for a patient of equivalent starting weight at the equivalent point in treatment — STEP 1 for semaglutide and SURMOUNT-1 for tirzepatide.

Cumulative medication exposure was calculated from logged dose history and expressed as a percentage of the exposure expected under the standard escalation protocol over the same treatment duration. This is an observational, indirect comparison against published trial data, not a randomized or controlled analysis.

Semaglutide (n = 30)

Mean treatment duration 30.8 weeks. Weight loss rate: 1.18 lb/week (1.3× the STEP 1 predicted rate of 0.74 lb/week).[13] Average weekly exposure: 0.71 mg — approximately 51% of the standard 1.40 mg/week. Eighty percent outperformed their individual trial-predicted trajectory.

Tirzepatide (n = 82)

Mean treatment duration 22.7 weeks. Weight loss rate: 1.21 lb/week (1.1× the SURMOUNT-1 predicted rate of 0.96 lb/week).[14] Average weekly exposure: 4.05 mg — approximately 53% of the standard 7.70 mg/week. Ninety-one percent achieved at least 5% body weight loss.

Combined Signal

Across both medications: 92% achieved ≥5% weight loss, 75% achieved ≥10%. Zero patients discontinued due to medication-induced side effects — compared to 6–7% in

the trials themselves[13,18] and 46–65% in real-world cohort studies[1–3]. Discontinuation was defined as permanent cessation of therapy attributed by the patient or provider to intolerable medication-related adverse effects. Patients who discontinued for other reasons — including cost, insurance changes, personal preference, goal weight achievement, or loss to follow-up — were not counted in this metric.

The most clinically significant finding may not be the weight loss rate. It is the discontinuation figure. The medications are unchanged. What changed was the dosing framework — and with it, the experience of treatment.

Limitations

These are observational data from a single concierge telehealth practice. The cohort was not randomized or blinded. Selection effects may influence results, and the concierge telehealth model includes weekly asynchronous clinician engagement and lifestyle guidance that differ from trial conditions and may independently contribute to outcomes. The comparison against trial predictions is not equivalent to a controlled comparison against a standard-of-care arm. Body composition data are not available for the full cohort. These results are presented as hypothesis-generating rather than confirmatory.

However, the theoretical basis does not depend on this cohort: pharmacokinetic-guided dosing is established practice for immunosuppressants, lithium, and antimicrobial agents.[15,19] MyLevel applies the same principles to medications that accumulate predictably and produce highly variable individual responses.[17] The cohort demonstrates that the theory translates into practice — producing a consistent signal across two pharmacokinetically distinct agents in a population broader and more complex than registration trial populations.

VII. CLINICIAN OVERSIGHT AND SAFETY

Standard protocols produce continuous, invisible, unstructured medication level increases interspersed with large dose escalations.[10,12] MyLevel replaces this with small, precise, visible adjustments. It functions as a telemetry system for GLP-1 therapy — providing continuous visibility into medication exposure where previously there was none.

Patients already frequently modify their own dosing outside the clinical encounter — splitting doses, delaying injections, skipping doses, or following advice from outside the clinical relationship.[2,21] MyLevel channels this existing patient agency into a visible, clinician-governed system.

Once a therapeutic range is identified, most patients require only occasional adjustments. The dashboard provides asynchronous review of dosing, levels, weight, and symptoms. The frequency and severity of side effects tends to be far lower than under standard dosing,[8] reducing clinical workload. The titration process functions more like a continuously variable transmission than a gearbox with only four settings.

The platform supports two modes: approval-required (all changes need clinician authorization) and supervised self-titration (patient adjusts within clinician-defined boundaries, with all changes logged and visible). A warning triggers if a patient attempts to adjust more than two units in a single cycle — for context, a single standard protocol dose increase is equivalent to nearly 10 MyLevel increases.[13–14]

The Disappearing Patient Problem

Within standard prescribing, a patient who develops side effects and stops taking medication may simply vanish.[1–3] Within MyLevel, a patient who stops dosing produces an immediate visible signal — the exposure curve flatlines, the dashboard flags the gap, and the care team can initiate outreach while there is still time to act: *is there something we can adjust?*

VIII. IMPLICATIONS AND CONCLUSIONS

GLP-1 medications represent a rapidly growing component of pharmacy spend.[20] If these findings generalize, the implications for payers are substantial: comparable or superior outcomes at approximately half the medication consumption, with a higher proportion of patients completing treatment. Current discontinuation rates of 46–64%[1–3] represent a large population who stopped treatment because the dosing model pushed them past their tolerable range. A framework that keeps more patients in treatment at lower per-patient consumption improves both outcomes and economics simultaneously.

GLP-1 medications accumulate predictably with repeated dosing.[10–12] Their clinical effects are driven by medication level.[17] Individual patients have meaningfully different therapeutic windows.[16,17] The standard dose-escalation protocol has no reliable mechanism to find or hold those windows.[13–14] These facts follow directly from the published pharmacokinetics of the drug class.

MyLevel applies established pharmacokinetic principles to close the gap between what these medications make possible and what the current dosing infrastructure delivers. For clinicians, the standard protocol will continue to work for patients with wide therapeutic windows and robust tolerability. For the substantial proportion who do not fit that profile,[1–3,21] a more precise framework is available.

The medications are ready. The pharmacology is understood.[10–12] The tools now exist. The molecules themselves are the therapeutic engine. Precision dosing is the transmission that allows that engine to run at the rpm it was designed for. It is time to move toward truly individualized, level-based dosing that improves outcomes, decreases attrition, and realizes the full promise of these extraordinary medications.

IX. DISCLOSURES

IJE is the founder and CEO of MyLevel, Inc. and the named inventor on U.S. patent application No. 63/949,189. Voafit is operated by the author. The observational data were collected within the author's clinical practice. These relationships represent potential conflicts of interest. No external funding was received. The MyLevel dosing platform is available for licensing to clinics and health systems seeking to implement exposure-guided dosing.

X. REFERENCES

1. Do D, Lee T, Peasah SK, Good CB, Inneh A, Patel U. GLP-1 receptor agonist discontinuation among patients with obesity and/or type 2 diabetes. *JAMA Netw Open*. 2024;7(5):e2413172.

2. Gleason PP, Urick BY, Marshall LZ, Friedlander N, Qiu Y, Leslie RS. Real-world persistence and adherence to glucagon-like peptide-1 receptor agonists among obese commercially insured adults without diabetes. *J Manag Care Spec Pharm*. 2024;30(6):583–598.

3. Khan SS, Ndumele CE, Kazi DS. Discontinuation of glucagon-like peptide-1 receptor agonists. *JAMA*. 2025;333(2):113–114.

4. Wilding JPH, Batterham RL, Davies M, et al. Impact of semaglutide on body composition in adults with overweight or obesity: exploratory analysis of the STEP 1 study. *J Endocr Soc*. 2021;5(Suppl 1):A16–A17.

5. Conte C, Hall KD, Klein S. Is weight loss–induced muscle mass loss clinically relevant? *JAMA*. 2024;332(1):9–10.

6. Jensen SBK, Lundgren JR, Janus C, et al. Bone health after exercise alone, GLP-1 receptor agonist treatment, or combination treatment: a secondary analysis of a randomized clinical trial. *JAMA Netw Open*. 2024;7(6):e2416775.

7. Hansen MS, Wölfel EM, Jeromdesella S, et al. Once-weekly semaglutide versus placebo in adults with increased fracture risk: a randomised, double-blinded, two-centre, phase 2 trial. *EClinicalMedicine*. 2024;72:102624.

8. Wharton S, Calanna S, Davies M, et al. Gastrointestinal tolerability of once-weekly semaglutide 2.4 mg in adults with overweight or obesity, and the relationship between gastrointestinal adverse events and weight loss. *Diabetes Obes Metab*. 2022;24(1):94–105.

9. Bezin J, Gouverneur A, Pénichon M, et al. Gastrointestinal adverse events associated with GLP-1 RA in non-diabetic patients with overweight or obesity: a systematic review and network meta-analysis. *Int J Obes*. 2025. doi:10.1038/s41366-025-01859-6.

10. Hall S, Isaacs D, Clements JN. Pharmacokinetics and clinical implications of semaglutide: a new glucagon-like peptide (GLP)-1 receptor agonist. *Clin Pharmacokinet*. 2018;57(12):1529–1538.

11. Coskun T, Urva S, Roell WC, et al. Population pharmacokinetics of the GIP/GLP-1 receptor agonist tirzepatide. *Clin Pharmacol Ther*. 2024;115(4):736–747.

12. Overgaard RV, Delff PH, Petri KCC, Anderson TW, Flint A, Ingwersen SH. Population pharmacokinetics of semaglutide for type 2 diabetes. *Diabetes Ther*. 2019;10(2):649–662.

13. Wilding JPH, Batterham RL, Calanna S, et al. Once-weekly semaglutide in adults with overweight or obesity. *N Engl J Med*. 2021;384(11):989–1002.

14. Jastreboff AM, Aronne LJ, Ahmad NN, et al. Tirzepatide once weekly for the treatment of obesity. *N Engl J Med*. 2022;387(3):205–216.

15. Rybak MJ, Le J, Lodise TP, et al. Therapeutic monitoring of vancomycin for serious methicillin-resistant *Staphylococcus aureus* infections: a revised consensus guideline and review. *Am J Health Syst Pharm*. 2020;77(11):835–864.

16. Brunton LL, Hilal-Dandan R, Knollmann BC, eds. *Goodman & Gilman's: The Pharmacological Basis of Therapeutics*. 13th ed. McGraw-Hill; 2018.

17. Overgaard RV, Hertz CL, Ingwersen SH, Navarria A, Drucker DJ. Levels of circulating semaglutide determine reductions in HbA1c and body weight in people with type 2 diabetes. *Cell Rep Med*. 2021;2(9):100387.

18. Wadden TA, Bailey TS, Billings LK, et al. Effect of subcutaneous semaglutide vs placebo as an adjunct to intensive behavioral therapy on body weight in adults with overweight or obesity: the STEP 3 randomized clinical trial. *JAMA*. 2021;325(14):1403–1413.

19. Kahan BD. Therapeutic drug monitoring of cyclosporine: 20 years of progress. *Transplant Proc*. 2004;36(2 Suppl):378S–391S.

20. Tichy EM, Hoffman JM, Suda KJ, et al. National trends in prescription drug expenditures and projections for 2025. *Am J Health Syst Pharm*. 2025;82(11):659–677.

21. Tsilidis KK, Patorno E, Engel PA, et al. Real-world evidence on the utilization, clinical and comparative effectiveness, and adverse effects of newer GLP-1RA–based weight-loss therapies. *Diabetes Obes Metab*. 2025;27(4):1711–1740.