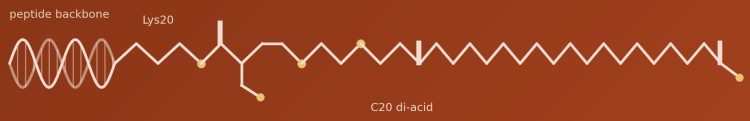
A CLINICAL MONOGRAPH · CONTINUING MEDICAL EDUCATION




# Tirzepatide

The First Dual GIP / GLP-1 Receptor Agonist — Mechanism of Action, Whole-Body Effects, Patient Selection, Dosing and Safety

## A complete mechanism-to-bedside reference

How tirzepatide works at two receptors — augmenting glucose-dependent insulin secretion, suppressing glucagon-driven gluconeogenesis, and reversing insulin resistance through both GLP-1 and a distinctive GIP adipose-tissue arm — and what that mechanism delivers across glycaemia, weight, cardiovascular, renal, hepatic, sleep-apnoea and addiction outcomes. Followed by patient selection, full dosing (including micro- and macro-dosing), and a mechanism-based account of every major adverse effect: why it happens, who is at risk, how to prevent it, and when it is not truly the drug.

For continuing medical education. June 2026 edition. Trial figures and regulatory landscape current to mid-2026.

# Contents


1 From the Incretins to a Dual Agonist

2 Molecular Structure and Biased Agonism

3 Mechanism I — Receptor Signalling and Glucose-Dependent Insulin Secretion

4 Mechanism II — Glucagon Suppression and the Control of Gluconeogenesis

5 Mechanism III — Insulin Resistance, the GIP Adipose Arm and Weight Loss

6 Mechanism IV — Appetite, Gastric Emptying and the Integrated Phenotype

7 Mechanism V — Reward, Cravings and Addiction: Food, Alcohol and Nicotine

8 Pharmacokinetics

9 Patient Selection — For Whom, and For Whom Not

10 Dosing I — Subcutaneous Titration

11 Dosing II — Microdosing, Macrodosing and the Current Landscape

12 Glycaemic and Weight Efficacy — SURPASS and SURMOUNT

13 Cardiovascular Outcomes and Their Mechanism — SURPASS-CVOT and SUMMIT

14 Renal Outcomes and Their Mechanism — SURPASS-4

15 Liver, Heart Failure and Sleep Apnoea — SYNERGY-NASH, SUMMIT, SURMOUNT-OSA

16 The Eye — Diabetic Retinopathy and the Optic-Neuropathy Question

17 Adverse Effects — Mechanisms, Predictors, Prevention and Confounders

18 Practical Management — Diet, Muscle, Plateau and Discontinuation

19 The Incretin Landscape and Future Directions (2026)

R References


**How to read this monograph**

This edition is built to make the reader fluent in tirzepatide from molecule to bedside. Chapters 3–7 develop the pharmacology in depth — including the distinctive **GIP** contribution to insulin sensitivity and tolerability, and the reward-circuit basis of reduced food, alcohol and nicotine craving. Chapters 12–16 then show what that mechanism produces in hard outcomes and explain **why**. Chapter 17 treats safety mechanistically. **PEARL** boxes flag practice points; each chapter closes with a **Key Points** box. Numeric superscripts refer to the reference list.

01 ORIGINS

# From the Incretins to a Dual Agonist

Tirzepatide occupies a distinctive place in the history of metabolic medicine: it is the first agent to validate, in humans and at scale, the idea that engaging *two* incretin receptors with a single molecule can outperform engaging one. To understand why this matters — and why a hormone, GIP, that had been written off as a therapeutic dead end turned out to be a powerful partner — one must first understand the incretin system in depth. Almost every clinically relevant property of tirzepatide, from its unprecedented degree of weight loss to its gastrointestinal profile, its glucose-dependent safety, its effects on adipose tissue and its dosing schedule, follows from the physiology of the two hormones it imitates and from the engineering that turned those fragile peptides into a once-weekly drug. This opening chapter develops that foundation: the discovery of the incretin effect; the two incretin hormones GIP and GLP-1 and how they are made and released; their full repertoire of actions; the phenomenon of GIP resistance that long discouraged GIP-based therapy and the way tirzepatide overcame it; the pharmacological obstacle that kept the native peptides from ever becoming drugs; and the engineering arc that produced a balanced, long-acting dual agonist. The remaining chapters build on this, but the logic of the whole monograph is set here.

## 1.1 The incretin effect — a century-old observation

The conceptual seed was planted early in the twentieth century. In 1902 Bayliss and Starling, having just coined the word **hormone**, speculated that the intestinal mucosa might release a blood-borne factor capable of stimulating the pancreas. In 1932 the Belgian physiologist Jean La Barre proposed the term **incretin** — from "INtestinal seCRETion of INsulin" — for a putative gut substance that lowered blood glucose by promoting endocrine pancreatic secretion. For decades the idea remained unproven, because the tools to measure circulating insulin did not exist.

That changed with the radioimmunoassay for insulin developed by Berson and Yalow in the late 1950s. Within a few years two landmark experiments established the phenomenon beyond doubt. In 1964, working independently, Elrick and colleagues and McIntyre and colleagues showed that an *oral* glucose load produced substantially more insulin than an *intravenous* glucose infusion engineered to reproduce the same blood-glucose excursion. Perley and Kipnis confirmed and extended these findings in 1967. The conclusion was inescapable: glucose arriving through the gut elicits more insulin than the same glucose delivered directly into the bloodstream, and the difference must be mediated by signals released from the intestine during nutrient absorption. That difference is the **incretin effect**.

Quantitatively, the incretin effect is a dominant determinant of post-prandial insulin secretion: in healthy individuals the gut-derived incretin signals are responsible for roughly 50–70% of the total insulin secreted in response to an oral glucose load, with the proportional contribution rising as the glucose load increases. For most of a normal day's meals, then, the majority of the insulin response is driven not by glucose acting directly on the beta cell but by hormones the gut releases during nutrient absorption. This reframes the beta cell as the downstream executor of a gut-initiated programme — and it explains why drugs that mimic those gut signals can exert such large and physiological effects on glucose handling.

## 1.2 Two incretin hormones — GIP and GLP-1

Two molecules account for essentially the entire incretin effect in humans, and tirzepatide is unique among approved agents in targeting both.

The first to be characterised was **glucose-dependent insulinotropic polypeptide** (GIP), originally named gastric inhibitory polypeptide for an early-described ability to inhibit gastric acid secretion at supraphysiological doses. GIP is a 42-amino-acid peptide secreted by the enteroendocrine **K cells**, concentrated in the proximal small intestine — the duodenum and proximal jejunum. As nutrients, particularly fat and glucose, reach this region, K cells release GIP, which acts on the GIP receptor on the pancreatic beta cell to augment glucose-stimulated insulin secretion. In health, GIP is the quantitatively larger contributor to the incretin effect. Its receptor, however, is expressed well beyond the beta cell — notably on adipocytes, where GIP influences lipid storage, blood flow and insulin sensitivity, and in the central nervous system, where it participates in the control of food intake. This broad distribution becomes central to tirzepatide's mechanism (Chapter 5) and is the reason a GIP-engaging drug does more than simply add to insulin secretion.

The second incretin is **glucagon-like peptide-1** (GLP-1), produced by the enteroendocrine **L cells**, distributed predominantly in the distal small intestine and colon. GLP-1 derives from the **proglucagon** gene, whose primary transcript is processed in a tissue-specific manner: in the pancreatic alpha cell, prohormone convertase 2 yields glucagon; in the intestinal L cell and in certain brainstem neurons, prohormone convertase 1/3 yields GLP-1, GLP-2, oxyntomodulin and glicentin. One gene, several biologically distinct peptides, determined by which convertase a cell expresses. This shared proglucagon ancestry also explains the considerable structural homology between GLP-1, GIP and glucagon — homology that the engineers of multi-agonist drugs exploit, since a single peptide can be shaped to engage more than one of these related receptors. Native GLP-1 circulates as GLP-1(7-37) and the more abundant amidated GLP-1(7-36)amide, both high-affinity agonists at the GLP-1 receptor.

## 1.3 How the incretins are secreted

Both K cells and L cells are open-type enteroendocrine cells, with apical microvilli sensing the gut lumen and basolateral surfaces secreting into the capillary network. They respond to all three macronutrient classes. Carbohydrate, particularly glucose, is sensed partly through the sodium-glucose cotransporter SGLT1, whose electrogenic transport depolarises the cell. Fat — especially the monoacylglycerols and long-chain fatty acids produced by triglyceride digestion — is a potent stimulus acting through G-protein-coupled receptors including GPR119, GPR40 and GPR120; fat is an especially strong trigger for GIP from proximal K cells. Protein and its digestion products also stimulate secretion, and bile acids acting on TGR5 contribute to GLP-1 release. As with GLP-1, GIP appears in the circulation within minutes of eating, reflecting both rapid proximal nutrient contact and neuro-hormonal anticipatory signalling. The system is thus tuned to nutrient ingestion — and because the insulinotropic actions of both hormones are glucose-dependent, they amplify insulin secretion when glucose is high and fall silent when it is normal or low.

02 CHEMISTRY

# Molecular Structure and Biased Agonism

If Chapter 1 explained why a dual incretin agonist was worth building, this chapter explains how a single peptide was engineered to engage two different receptors at once, to last a week in the circulation, and to signal in a deliberately tuned, "biased" manner that contributes to both its efficacy and its tolerability. Tirzepatide is among the most sophisticated peptide drugs in clinical use, and almost every clinical property a prescriber relies on — its dual mechanism, its once-weekly dosing, its glucose-dependent safety, the absence (so far) of an oral form, and even the way its tolerability differs from a pure GLP-1 agonist — follows from the structural and pharmacological choices described here. This chapter develops them in the depth those clinical inferences require.

## 2.1 The design brief — one peptide, two receptors

The central design challenge of tirzepatide was unprecedented among approved incretin drugs: to create a *single* molecule that activates both the GIP receptor and the GLP-1 receptor with the right balance of potency at each, while also resisting enzymatic degradation and binding albumin for once-weekly dosing. This is far harder than designing a selective agonist. A molecule must be shaped to fit two distinct (though related) receptor binding sites, and the relative strength of its action at each must be tuned, because the two receptors contribute different and not simply additive effects (Chapter 1). The solution was to start from one of the two native incretin sequences and modify it until it engaged both receptors appropriately — a strategy of building outward from a single hormonal scaffold rather than fusing two separate agonists together.

## 2.2 The GIP-based backbone

Tirzepatide is built upon the sequence of **GIP**, the 42-amino-acid glucose-dependent insulinotropic polypeptide, rather than upon GLP-1. This is a notable and deliberate choice. The designers took the GIP backbone — which natively activates the GIP receptor — and introduced modifications that conferred, in addition, potent activation of the GLP-1 receptor, exploiting the structural homology between the two class B receptors and their ligands (Chapter 1). The finished molecule is a 39-amino-acid synthetic peptide that is recognisably GIP-derived but engineered to be a dual agonist. Building on GIP, rather than on GLP-1, reflects both the homology that makes cross-activation achievable and the goal of retaining genuine GIP-receptor engagement, with its distinctive adipose and central actions, as a core part of the molecule's identity rather than an afterthought.

As with semaglutide, the design philosophy was to make targeted, purposeful modifications: substitutions to confer protease resistance and to tune dual-receptor activity, and an acylation to confer the albumin binding required for weekly dosing. The result is a single, well-defined chemical entity whose every structural feature serves the dual-agonist, long-acting brief.

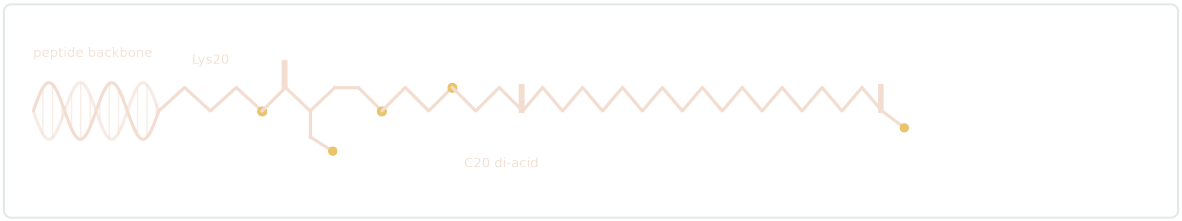



**Figure 2.1.** Schematic of tirzepatide: a GIP-based peptide backbone (helical ribbon) bearing, at a lysine, the acyl side chain — a γ-glutamate spacer and OEG units (oxygen atoms marked) linking the peptide to a C20 fatty di-acid that binds albumin.

## 2.3 The Aib substitutions — protease resistance and dual activity

Like GLP-1, native GIP is a substrate for **DPP-4**, which cleaves its N-terminal dipeptide and inactivates it within minutes (Chapter 1). The structural solution mirrors that used in semaglutide but is applied at the corresponding GIP position: substitution with **α-aminoisobutyric acid** (Aib), the sterically hindered, doubly methylated amino acid that obstructs the DPP-4 active site and simultaneously stabilises the helical conformation important for receptor activation. Tirzepatide incorporates Aib at two positions in its sequence (commonly described as positions 2 and 13 in the GIP numbering). The N-terminal Aib confers DPP-4 resistance, as in semaglutide. The additional Aib substitution contributes to the molecule's conformational stability and, importantly, to the fine-tuning of its activity across the two receptors. These substitutions are part of how the designers achieved both durability and the desired balance of dual-receptor engagement from a single backbone.

## 2.4 The acylation — a C20 fatty diacid for once-weekly dosing

To achieve the once-weekly half-life, tirzepatide uses the same **acylation principle** as the liraglutide–semaglutide lineage, but with a longer fatty acid. A **C20 fatty diacid** (a twenty-carbon chain with a carboxylic acid at each end) is attached to the side chain of a lysine in the peptide through a linker built from a γ-glutamate (γGlu) spacer and OEG (8-amino-3,6-dioxaoctanoic acid) units — the same modular linker chemistry seen in semaglutide. Reading outward from the peptide: lysine side chain → γGlu → OEG spacer(s) → C20 di-acid.

The mechanism of protraction is identical in principle to that described for semaglutide. The fatty-acid chain inserts into a hydrophobic binding pocket of circulating albumin, so that the great majority of the drug travels albumin-bound — protected from DPP-4 and from other proteases, shielded from glomerular filtration by the size of the albumin complex, and released slowly from this depot to maintain a steady free fraction at the receptors. The C20 di-acid binds albumin slightly more avidly than semaglutide's C18, contributing to tirzepatide's prolonged residence. The resulting half-life of approximately five days supports once-weekly dosing and produces the smooth, sustained dual-receptor stimulation that the slow titration of Chapter 10 then ramps up gradually. As with semaglutide, this depot has direct clinical consequences: a missed weekly dose can be taken several days late, and the drug's effects — and adverse effects — persist for weeks after the last dose, which matters for pregnancy planning and for managing serious adverse events.

03 MECHANISM I

# Receptor Signalling and Glucose-Dependent Insulin Secretion

This is the first of the five mechanism chapters, and it establishes the property that defines the safety of the whole incretin class: that tirzepatide enhances insulin secretion *only when glucose is elevated*. For a dual agonist there is an added dimension — two incretin receptors, GIP and GLP-1, both converge on the same intracellular second messenger in the beta cell, so the chapter must also explain what the GIP arm adds and why the long dismissal of GIP at the beta cell was only half the story. A clinician who grasps this mechanism can explain why tirzepatide lowers glucose powerfully yet rarely causes hypoglycaemia on its own, why hypoglycaemia (when it occurs) is the fault of a co-prescribed sulfonylurea or insulin, why the drug works best in patients with residual beta-cell function, and why combining two incretin signals amplifies the effect. To get there, the chapter first reconstructs how the healthy beta cell senses glucose, then shows where GIP- and GLP-1-receptor signalling intervene and why that intervention is gated by glucose.

## 3.1 The beta cell as a glucose sensor

Insulin secretion translates the ambient glucose concentration into a proportional amount of secreted insulin, through what is called the **triggering pathway**. Glucose enters the beta cell through facilitative transporters down its gradient, so intracellular glucose tracks the blood level. It is then phosphorylated by **glucokinase**, the true "glucose sensor": glucokinase has a high Michaelis constant, so its activity rises steeply across the physiological glucose range and sets the threshold and sensitivity of secretion. Phosphorylated glucose enters glycolysis and mitochondrial oxidation, raising the intracellular **ATP-to-ADP ratio** in proportion to the glucose load.

This rising ATP/ADP ratio closes the **ATP-sensitive potassium channel** ($K_{ATP}$), an octamer of four pore-forming Kir6.2 subunits and four regulatory SUR1 (sulfonylurea receptor) subunits. At low glucose the channels are open, holding the membrane hyperpolarised at rest; as ATP rises it binds Kir6.2 and closes them, and the membrane **depolarises**. The $K_{ATP}$ channel is the master switch converting the metabolic glucose signal into a voltage change — and it is the direct target of sulfonylureas, which close it through SUR1 regardless of glucose. Depolarisation opens voltage-gated calcium channels; calcium floods in; and the rise in intracellular calcium triggers the **exocytosis** of insulin granules. The chain — glucose → glucokinase → ATP/ADP rise → $K_{ATP}$ closure → depolarisation → calcium influx → insulin exocytosis — is proportional to glucose at every step, so insulin output tracks the glucose level. Secretion is biphasic: a rapid **first phase** from pre-docked granules, then a sustained **second phase**; loss of first-phase secretion is among the earliest beta-cell defects in type 2 diabetes.

## 3.2 From the incretin receptors to cyclic AMP — convergence of two receptors

Here the dual agonist differs instructively from a selective GLP-1 agonist. Both the **GLP-1 receptor** and the **GIP receptor** are class B G-protein-coupled receptors expressed on the beta cell, and **both couple to the stimulatory G-protein $G_s$**. When tirzepatide activates either receptor, $G_s$ activates **adenylyl cyclase**, which synthesises the second messenger **cyclic AMP (cAMP)** from ATP. The two incretin signals therefore *converge* on the same intracellular pathway: a rise in cAMP. This convergence is the molecular reason a dual agonist can amplify insulin secretion more than a selective agonist — two receptor inputs feed the same amplifying second messenger. It is also why the glucose-dependent safety property is preserved despite the addition of a second receptor: both inputs raise cAMP, and cAMP, as the next sections show, can only amplify a glucose-initiated process.

## 3.3 The two arms of cAMP action — PKA and Epac2

Cyclic AMP acts through two principal effectors. **Protein kinase A** (PKA) is the classical cAMP-dependent kinase: cAMP frees its catalytic subunits to phosphorylate ion channels, components of the exocytotic machinery and transcription factors. **Epac2** (exchange protein directly activated by cAMP 2) acts independently of PKA, principally on the secretory machinery and intracellular calcium handling. The existence of two parallel effectors — a broadly acting kinase and a guanine-nucleotide exchange factor focused on exocytosis and calcium — is why incretin signalling amplifies several distinct steps of insulin secretion rather than one. In the dual agonist, both GIP- and GLP-1-receptor activation feed cAMP into this same PKA/Epac2 machinery.

## 3.4 How incretin signalling amplifies each step

The cAMP/PKA/Epac2 signal does not create a new secretory pathway; it *amplifies the existing glucose-triggered one*, at multiple points. At the **$K_{ATP}$ channel**, PKA and Epac2 further promote closure and depolarisation. At the level of **calcium**, cAMP enhances entry through voltage-gated channels and mobilises calcium from the endoplasmic reticulum (Epac2 promoting calcium-induced calcium release through ryanodine receptors). At the **secretory machinery**, the signal acts on the proteins that dock, prime and fuse granules, enlarging the **readily-releasable pool** and improving the efficiency with which calcium drives exocytosis. The net effect is that, for any given glucose stimulus, more insulin is secreted. Table 3.1 maps the amplification onto each step.

**Table 3.1. Glucose-stimulated insulin secretion and where incretin signalling amplifies it.**

04 MECHANISM II

# Glucagon Suppression and the Control of Gluconeogenesis

Chapter 3 explained how tirzepatide increases insulin secretion when glucose is high. This chapter addresses the other side of glucose control — the suppression of the liver's own glucose production — which is at least as important clinically, because in type 2 diabetes a large part of the problem, especially the *fasting* hyperglycaemia that drives much of the HbA1c, originates in too much glucose pouring out of the liver around the clock. For a dual agonist there is an added subtlety: the two incretins act *oppositely* on glucagon — GLP-1 suppresses it, GIP can raise it at euglycaemia — so the chapter must explain not only how gluconeogenesis is driven and stopped but how the net glucagon effect of tirzepatide emerges from these competing influences. This is among the most mechanistically detailed chapters in the monograph, answering precisely how gluconeogenesis happens and how raising insulin stops it, at the level of the enzymes and transcription factors involved.

## 4.1 Hepatic glucose production and fasting hyperglycaemia

The liver maintains blood glucose between meals and overnight by releasing glucose into the circulation — essential for the brain, which needs a continuous supply. It does this by **glycogenolysis** (breakdown of stored glycogen, dominant in early fasting but limited by store size) and by **gluconeogenesis** (synthesis of new glucose from non-carbohydrate precursors, sustaining glucose production through prolonged fasting). In health this is exquisitely regulated — switched on in fasting, off after meals when insulin rises. In type 2 diabetes the regulation fails: hepatic glucose production becomes excessive and is no longer suppressed in the fed state, so the liver pours out glucose even when blood glucose is already high. This *unrestrained hepatic glucose output*, driven predominantly by excessive gluconeogenesis, is the principal source of fasting hyperglycaemia in type 2 diabetes. Reining it in addresses a core mechanism of the disease.

## 4.2 The biochemistry of gluconeogenesis

Gluconeogenesis builds glucose from **lactate** (from glycolysis, via the Cori cycle), **glycerol** (from adipose lipolysis) and **glucogenic amino acids**, especially alanine (from muscle, via the glucose–alanine cycle). It is essentially glycolysis in reverse, but three irreversible glycolytic steps must be bypassed by distinct enzymes — the regulated control points that glucagon induces and insulin represses. The first bypass converts pyruvate to phosphoenolpyruvate via **pyruvate carboxylase** (yielding oxaloacetate) and **phosphoenolpyruvate carboxykinase** (**PEPCK**, *PCK1*). The second is **fructose-1,6-bisphosphatase** (FBPase-1), the reciprocal of glycolytic PFK-1. The third is **glucose-6-phosphatase** (**G6Pase**, *G6PC*), which releases free glucose. PEPCK and G6Pase are the enzymes whose transcription is the principal long-term lever of hormonal control.

A crucial regulatory node is the signalling metabolite **fructose-2,6-bisphosphate**: it activates glycolysis (stimulating PFK-1) and inhibits gluconeogenesis (inhibiting FBPase-1). Its level, set by a phosphorylation-controlled bifunctional enzyme, acts as a master switch — high

fructose-2,6-bisphosphate means the liver burns glucose; low means it makes glucose. Glucagon lowers it; insulin raises it; and that is part of how the two hormones flip the liver between glucose production and consumption.

## 4.3 How glucagon drives gluconeogenesis

Glucagon binds the hepatocyte **glucagon receptor** (a $G_s$-coupled class B GPCR related to the GLP-1 and GIP receptors), raising cAMP and activating **PKA**, which drives hepatic glucose output through three coordinated actions (Table 4.1). It stimulates **glycogenolysis** (via phosphorylase kinase → glycogen phosphorylase, while inhibiting glycogen synthase). It **induces the gluconeogenic genes** by activating CREB, which with PGC-1α and FOXO1 drives transcription of PEPCK and G6Pase. And it **redirects flux** by lowering fructose-2,6-bisphosphate, relieving FBPase-1 inhibition and removing PFK-1 activation, favouring gluconeogenesis over glycolysis.

**Table 4.1. The glucagon → cAMP → PKA programme that raises hepatic glucose output.**

| PKA action | Target | Effect |
|---|---|---|
| Activates glycogen breakdown | Phosphorylase kinase → glycogen phosphorylase | ↑ glycogenolysis |
| Induces gluconeogenic genes | CREB → PGC-1α/FOXO1 → **PEPCK (PCK1)** and **glucose-6-phosphatase (G6PC)** | ↑ gluconeogenic capacity |
| Redirects flux to glucose | ↓ fructose-2,6-bisphosphate → relieves FBPase-1 inhibition, inhibits PFK-1 | Favours gluconeogenesis over glycolysis |

## 4.4 The insulin-to-glucagon ratio — the master switch

The liver responds not to the absolute level of either hormone but to the **portal insulin-to-glucagon ratio**. High ratio (after a meal) stops glucose production and stores glycogen; low ratio (fasting) makes and releases glucose. Glucagon and insulin act on overlapping targets in opposite directions, so a therapy can normalise hepatic glucose output by lowering glucagon, raising insulin action, or both. The dual agonist's net effect on this ratio is shaped by the competing glucagon influences of its two arms, examined below.

## 4.5 The diabetic defect — hyperglucagonaemia and hepatic insulin resistance

In type 2 diabetes the ratio is deranged in both terms: glucagon secretion is inappropriately high (**hyperglucagonaemia**, with failure to suppress after meals), and the liver is *resistant* to insulin, so the effective insulin signal is reduced. An elevated glucagon signal drives the PEPCK/G6Pase programme while a weakened insulin signal fails to restrain it — producing excessive, poorly suppressed gluconeogenesis and fasting hyperglycaemia. Correcting it requires shifting the ratio back.

05 MECHANISM III

# Insulin Resistance, the GIP Adipose Arm and Weight Loss

The two preceding chapters concerned insulin *secretion* and glucagon. But the defining lesion of type 2 diabetes, and the central problem in obesity-related metabolic disease, is insulin *resistance*. This chapter explains what insulin resistance is molecularly, why ectopic fat is its dominant cause, how inflammation contributes, and how weight loss reverses it — and then develops the feature that most distinguishes tirzepatide from a selective GLP-1 agonist: a *direct* action on adipose tissue through the GIP receptor that appears to improve insulin sensitivity beyond what the weight loss alone would predict. This GIP adipose arm is a leading explanation for why a dual agonist achieves greater metabolic benefit than GLP-1 agonism alone, and it is the mechanistic centre of this chapter.

## 5.1 What insulin does at the cell — the signalling cascade

Insulin binds the insulin receptor, a receptor tyrosine kinase, activating its kinase to phosphorylate the docking proteins **IRS-1 and IRS-2** on *tyrosine* residues. Tyrosine-phosphorylated IRS recruits and activates **PI3-kinase (PI3K)**, which generates $PIP_3$ and activates the central metabolic kinase **Akt**. In **muscle and adipose tissue**, Akt drives translocation of the glucose transporter **GLUT4** to the cell surface, admitting glucose — muscle being the largest site of insulin-stimulated glucose disposal. In the **liver**, Akt excludes FOXO1 from the nucleus (switching off gluconeogenesis, Chapter 4) and promotes glycogen synthesis. The IRS-1 → PI3K → Akt → GLUT4 axis is the pathway whose failure constitutes insulin resistance.

## 5.2 What insulin resistance is, molecularly

Insulin resistance is a defect in this cascade: cells receive the insulin signal but transmit it poorly, so a normal or elevated insulin level produces a subnormal response. The body compensates with hyperinsulinaemia until the beta cell can no longer keep pace. The key lesion is at IRS-1: instead of the productive *tyrosine* phosphorylation that propagates the signal, IRS-1 in the resistant state carries excessive inhibitory **serine phosphorylation**, which *uncouples* the receptor from PI3K/Akt. With IRS-1 uncoupled, PI3K and Akt are inadequately activated, GLUT4 translocation fails, and glucose uptake into muscle falls. The central question is what causes this inhibitory serine phosphorylation — and the answer, in most cases, is ectopic fat.

## 5.3 The dominant driver — ectopic fat and lipotoxicity

The best-characterised cause of insulin resistance is **ectopic fat** — lipid deposited where it does not belong — and its toxic effects, termed **lipotoxicity**. Adipose tissue is the designated fat store; when its capacity (particularly subcutaneous) is exceeded — the **adipose expandability / overflow** hypothesis — lipid deposits in muscle, liver and pancreas. It is this ectopic lipid, more than total fat, that drives resistance, which is why fat *distribution* matters more than weight alone.

The mechanistic chain is well defined. Excess intracellular lipid raises **diacylglycerol (DAG)**, a physiological activator of the **novel protein kinase C** isoforms — PKCθ in skeletal muscle and PKCε in liver. Activated by accumulating DAG, these kinases phosphorylate IRS-1 (and the insulin receptor) on inhibitory *serine* residues, uncoupling IRS-1 from the receptor, so PI3K/Akt signalling fails — in muscle reducing GLUT4 glucose uptake, in liver leaving gluconeogenesis unrestrained (Chapter 4). A parallel class of lipids, the **ceramides**, impairs Akt activation by related mechanisms. Ectopic fat is thus not inert storage but an active source of signalling lipids that switch off insulin signalling at its most vulnerable node.

**PEARL** Insulin resistance is, to a first approximation, a lipid problem masquerading as a glucose problem: ectopic fat → diacylglycerol → novel PKC → serine-phosphorylation of IRS-1 → uncoupling of the insulin receptor → failed GLUT4 glucose uptake. Every step is reversible by depleting the ectopic fat — and, as Section 5.5 argues, tirzepatide may attack the problem from a second direction by improving how adipose tissue stores fat in the first place.

## 5.4 Inflammation and glucotoxicity — parallel routes

Two further drivers reinforce lipotoxicity. Expanded, dysfunctional (especially visceral) adipose tissue recruits macrophages and secretes pro-inflammatory cytokines — **TNF-α** and **IL-6** — which activate the inflammatory kinases **JNK** and **IKKβ**; these, like the novel PKCs, serine-phosphorylate IRS-1, converging on the same lesion and raising **CRP**. Separately, **glucotoxicity** — the direct impairment of insulin secretion and action by chronic hyperglycaemia — creates a vicious cycle in which high glucose worsens both secretion and sensitivity. Lowering glucose and weight therefore relieves both the inflammatory and glucotoxic components.

## 5.5 The GIP-specific contribution — adipose tissue as a buffer

Here is where tirzepatide departs from a selective GLP-1 agonist. The GIP receptor is highly expressed on **adipocytes**, and the GIP arm of tirzepatide appears to act directly on adipose tissue to improve its core function as a metabolic *buffer*. The hypothesis, supported by preclinical and mechanistic human data, is that GIP-receptor activation enhances the ability of healthy adipose tissue to take up, esterify and safely store dietary lipid — promoting post-prandial blood flow to adipose tissue, supporting the trapping of fatty acids as triglyceride, and improving local adipose insulin sensitivity. If adipose tissue buffers dietary fat more effectively, less lipid overflows to ectopic sites; less ectopic lipid means less DAG, less novel-PKC activation, less serine phosphorylation of IRS-1, and therefore better insulin signalling in muscle and liver. In other words, where weight loss reverses insulin resistance by *removing* ectopic fat after the fact, the GIP adipose arm may *prevent* ectopic deposition by improving where fat is stored in the first place — attacking the same problem from the opposite end.

The clinical signature of this additional arm is that tirzepatide tends to improve insulin sensitivity *more than the weight loss alone would predict*. Studies of body composition and metabolic markers suggest that, for a given amount of weight lost, the dual agonist produces metabolic improvements — in insulin sensitivity and in some lipid and inflammatory parameters — that exceed those attributable to weight loss by itself, a difference reasonably attributed to the direct adipose actions of the GIP component. This is the mechanistic counterpart of the clinical observation (Chapter 12) that tirzepatide outperforms selective

06 MECHANISM IV

# Appetite, Gastric Emptying and the Integrated Phenotype

Weight loss — the engine of much of tirzepatide's metabolic benefit (Chapter 5) — is driven mainly by central and gastrointestinal actions: reduced hunger, earlier fullness, a quieting of preoccupation with food, and slowed gastric emptying. For a dual agonist there is an added dimension, because the GIP receptor, like the GLP-1 receptor, is expressed in the appetite and nausea circuitry of the brain, so the GIP arm contributes to both the satiety effect and, importantly, to tolerability. This chapter develops these appetite mechanisms, distinguishes homeostatic hunger (here) from hedonic craving (Chapter 7), and shows how the four mechanisms of Chapters 3–6 converge on a single cardiometabolic phenotype.

## 6.1 Two systems of appetite control

Food intake is governed by an intertwined pair of systems. The **homeostatic** system regulates energy balance and generates and relieves hunger, defending the body-weight set-point through gut and adipose hormones (Chapter 1). The **hedonic** (reward) system governs the desire for the pleasure of palatable food independent of need and is the seat of craving (Chapter 7). Both incretins act on both systems; this chapter concerns the homeostatic arm, while recognising that the two are anatomically interwoven.

## 6.2 The homeostatic hub — the arcuate nucleus and melanocortin system

The hub of homeostatic appetite control is the hypothalamic **arcuate nucleus**, with two opposing first-order neuron populations. **POMC/CART** neurons are *anorexigenic* — activating them suppresses appetite and promotes satiety (largely via α-MSH acting on melanocortin MC4 receptors). **AgRP/NPY** neurons are *orexigenic* — activating them powerfully stimulates hunger and antagonises the melanocortin system. Their balance, integrating signals of energy status (leptin, insulin, ghrelin, incretins), sets hunger moment to moment. GLP-1-receptor signalling shifts this balance toward satiety by *stimulating POMC/CART* and *inhibiting AgRP/NPY* neurons, reducing hunger and producing earlier fullness — a pharmacological resetting of the hunger–satiety balance rather than willpower, acting on the same melanocortin circuitry that defends the set-point.

## 6.3 The GIP central contribution — appetite and nausea

What distinguishes a dual agonist here is that **GIP receptors are also expressed centrally**, in regions governing food intake and in the nausea pathways of the hindbrain. The GIP arm therefore contributes to the central control of appetite alongside the GLP-1 arm — the two incretin signals converging on overlapping circuitry to reduce intake. More striking, and of real clinical value, is evidence that central **GIP-receptor signalling can attenuate nausea and emesis**, including the nausea driven by GLP-1-receptor activation. If GIP-receptor agonism dampens the emetic response while still contributing to satiety, then the GIP arm of tirzepatide may do double duty: adding to the appetite-suppressing effect while

simultaneously *improving the gastrointestinal tolerability* of the molecule. This is one of the leading explanations — alongside the imbalanced/biased pharmacology of Chapter 2 — for the clinical observation that tirzepatide can be escalated to very high efficacy with tolerability that, dose for effect, remains comparable to selective GLP-1 agonists despite the greater potency. The GIP arm, in other words, may help solve the central trade-off of incretin therapy, between efficacy and nausea, that limits selective agonists.

## 6.4 The hindbrain — area postrema, NTS, and how the drug reaches the brain

The **hindbrain** — the nucleus of the solitary tract (NTS) and the adjacent **area postrema** — is a second major appetite centre and the route by which a large, albumin-bound peripheral peptide exerts central effects. The area postrema is a **circumventricular organ** with an incomplete blood–brain barrier, allowing circulating peptides direct access; peripheral incretin signals also reach the NTS via **vagal afferents** from the gut. The area postrema is additionally the brain's principal **emetic sensor**, so the same accessibility that engages satiety circuitry underlies nausea (Chapter 17) — which is precisely the trade-off that the GIP arm (Section 6.3) appears to soften.

## 6.5 "Food noise" and delayed gastric emptying

Patients consistently report not only reduced hunger but the disappearance of **food noise** — the intrusive background preoccupation with food — which sits at the interface of homeostatic and hedonic systems (Chapter 7) and makes reduced intake feel effortless rather than restrictive. Separately, tirzepatide *slows gastric emptying* (predominantly a GLP-1-mediated effect), prolonging gastric distension (a satiety signal) and blunting post-prandial glucose by delivering nutrients more gradually. This effect shows **tachyphylaxis** — attenuating with continued exposure, so nausea and early satiety wane over weeks — but its persistence during escalation and with GI symptoms underlies the peri-operative aspiration caution of Chapter 17.

## 6.6 The integrated phenotype

The four mechanisms of Chapters 3–6 must be seen as a coordinated whole. They converge: augmented glucose-dependent insulin secretion at the **beta cell** (Chapter 3); suppression of glucagon-driven gluconeogenesis at the **liver** (Chapter 4); improved insulin sensitivity at **muscle and adipose** as ectopic fat falls, aided by the GIP adipose arm (Chapter 5); and reduced energy intake via appetite suppression and slowed emptying at the **brain and gut** (this chapter). These reinforce one another — reduced intake drives weight loss, improving insulin sensitivity, improving glycaemia, relieving glucotoxicity, improving beta-cell function — producing a coherent **cardiometabolic phenotype**: large reductions in HbA1c and weight with lower blood pressure, improved lipids, reduced hepatic fat and lower inflammatory markers. Figure 6.1 depicts this convergence.

07 MECHANISM V

# Reward, Cravings and Addiction: Food, Alcohol and Nicotine

Like other incretin agonists, tirzepatide quiets **craving** — patients describe the end of intrusive thoughts about food, and an emerging literature reports reduced desire for alcohol and tobacco. This is a mechanism distinct from gastric emptying or hypothalamic satiety, and the GLP-1 component of the molecule is its principal driver. For the dual agonist, two qualifications apply: the reward effect is attributed chiefly to GLP-1-receptor signalling (the role of GIP in reward is less defined), and the dedicated clinical evidence in substance-use disorders is at present thinner for tirzepatide specifically than for semaglutide, so much of the human data is best read as a class effect. This chapter develops the reward circuitry, explains how the drug dampens it, and reviews the evidence with appropriate care.

## 7.1 Why craving is a distinct mechanism

Everyday language conflates *hunger* and *craving*, but they are different. Hunger is the homeostatic drive to eat that arises from energy need, governed by the hypothalamic and gut mechanisms of Chapter 6. Craving is the hedonic pull toward a specific rewarding substance — a palatable food, a drink, a cigarette — that arises largely independent of energy need and is governed by the reward system. A full person can crave dessert; that is hedonic, not homeostatic. This matters because tirzepatide reduces both, through different circuitry, and the craving effect cannot be explained by slowed gastric emptying or POMC/AgRP satiety. Patients who report that food "lost its grip," or that a previously irresistible drink or cigarette lost its appeal, are describing a change in the reward system.

## 7.2 The mesolimbic reward circuit

The seat of reward and craving is the **mesolimbic dopamine system**. Dopamine neurons of the **ventral tegmental area** (VTA) project to the **nucleus accumbens** (NAc) and **prefrontal cortex** (PFC). When a reward — or a *cue* predicting one — is encountered, VTA neurons fire and release dopamine in the accumbens, assigning **incentive salience** to the cue: making it "wanted," grabbing attention and motivating consumption. Neuroscience distinguishes this dopamine-mediated "*wanting*" (the motivational pull) from opioid-mediated "*liking*" (consummatory pleasure); addiction is marked especially by excessive, cue-triggered *wanting*. The dopamine signal also encodes reward prediction, learning to fire to cues, which is how environmental triggers acquire their power to provoke craving. This circuit is the common final pathway for the reinforcing properties of food, alcohol, nicotine and other drugs — which is why one intervention acting on it can influence several behaviours.

## 7.3 GLP-1 receptors sit inside the reward system

GLP-1 receptors are expressed throughout the mesolimbic and mesocortical reward circuitry — VTA, nucleus accumbens, prefrontal cortex, lateral septum and the hindbrain NTS — and the brain contains its own GLP-1-producing neurons (the preproglucagon neurons of the NTS) projecting to these regions, so the reward system is natively wired to be modulated by GLP-1

signalling. GIP receptors are also present in the CNS and may modulate these circuits, but the anti-craving effect of tirzepatide is attributed chiefly to its GLP-1 arm; the contribution of the GIP arm to reward, if any, is not well defined. The drug therefore engages a physiological GLP-1 reward-regulating system rather than an incidental target.

## 7.4 How tirzepatide dampens the circuit

A long-acting peripheral agonist reaches these deep regions via **vagal afferents** to the NTS and via **circumventricular organs** (area postrema, median eminence) where the blood–brain barrier is incomplete (Chapter 6). Once engaged, GLP-1-receptor activation *inhibits VTA dopamine neurons* and reduces stimulated dopamine release in the nucleus accumbens; across preclinical studies the dopamine surge normally evoked by food, alcohol, nicotine and cocaine is blunted. The cue still occurs but now predicts a smaller dopamine signal, so its incentive salience — the "wanting" — shrinks. The drug also opposes **ghrelin**, which amplifies mesolimbic reward signalling, so tirzepatide shifts the reward balance away from craving from two directions, while modulating glutamatergic and GABAergic signalling within the circuit. The behavioural result is reduced craving and reduced consumption of reinforcing substances.

> **PEARL** "Food noise" disappearing is the same circuit change — reduced incentive salience of a reinforcing cue — that plausibly underlies reductions in alcohol and tobacco use. One mechanism, several behavioural read-outs, driven principally by the GLP-1 arm.

## 7.5 Food cravings and hedonic eating

Beyond the homeostatic satiety of Chapter 6, tirzepatide reduces the *reward value* of palatable, energy-dense food. In choice paradigms GLP-1-receptor agonism preferentially reduces intake of palatable food over bland chow — a hedonic signature. Clinically this is why patients lose interest disproportionately in snacks, sweets and "treat" foods, why binge-eating frequency falls, and why previously irresistible foods become a matter of indifference. Given that tirzepatide produces the largest weight loss of the licensed incretins, this reduction in food reward is an important contributor to the magnitude and durability of that loss, addressing not only how much patients are driven to eat but how strongly they are drawn to the most rewarding foods.

## 7.6 Alcohol use disorder

The strongest dedicated trial evidence in alcohol use disorder is for semaglutide — a randomised placebo-controlled trial (Hendershot, 2025) showing reduced consumption, craving and heavy-drinking days, the 26-week SEMALCO randomised trial, and large real-world pharmacoepidemiological cohorts. For *tirzepatide specifically*, the direct evidence is thinner but concordant: a study in individuals with obesity found that **both semaglutide and tirzepatide reduced alcohol consumption**, consistent with a class effect driven by the shared GLP-1-mediated reward mechanism. Because tirzepatide is a potent GLP-1-receptor agonist, the mechanistic expectation is that it should share the anti-alcohol effect, but dedicated randomised tirzepatide trials in alcohol use disorder are still needed to establish this directly. As with semaglutide, alcohol use disorder is *not* a licensed indication, and the drug should not be prescribed solely to treat it outside a trial — but for the common patient

08 PHARMACOKINETICS

# Absorption, Distribution, Metabolism and Elimination

With the mechanisms established, the monograph turns to the practical pharmacology that governs how tirzepatide is dosed and used. Pharmacokinetics is the direct source of the clinical rules that follow in Chapters 9 to 11: why titration is gradual, why a missed dose can be taken days late, why effects persist for weeks after stopping, why no renal dose adjustment is needed, why there are almost no metabolic drug interactions — and one tirzepatide-specific point with real clinical weight, the effect on oral contraceptive absorption. Every one of these rules is a kinetic consequence of the molecule designed in Chapter 2, and this chapter draws the lines explicitly.

**Table 8.1. Key pharmacokinetic parameters (subcutaneous).**

| Parameter | Value |
|---|---|
| Half-life | ≈ 5 days |
| Dosing interval | Once weekly |
| Time to steady state | ≈ 4 weeks |
| Bioavailability (SC) | ≈ 80% |
| Time to peak ($T_{max}$) | ~1–3 days |
| Plasma protein binding | > 99% (reversibly albumin-bound) |
| Albumin binding | C20 fatty-diacid acylation → prolonged half-life |
| Metabolism | Proteolysis + fatty-acid β-oxidation; no major CYP route |
| Renal/hepatic impairment | No dose adjustment required |
| Oral formulation | None approved (in development) |

## 8.1 Absorption — and the absence of an oral form

After subcutaneous injection, tirzepatide is absorbed slowly and largely completely, with an absolute bioavailability of approximately 80% and a time to peak plasma concentration of roughly one to three days. As with semaglutide, the slow absorption is a consequence of the albumin-binding design, which holds the drug in a depot rather than releasing it as a sharp peak, contributing to a smooth, low-variability concentration profile that reduces concentration-related side effects. Unlike semaglutide, tirzepatide has *no approved oral formulation*: the strategy that made oral semaglutide possible (the SNAC absorption enhancer, Chapter 2) has not been applied to an approved tirzepatide product, so the strict empty-stomach administration rules that apply to oral peptides are not relevant here. An oral tirzepatide is in development, but at present the drug is given only by subcutaneous injection, and absorption considerations reduce to the slow, near-complete subcutaneous uptake described here.

## 8.2 Distribution and the albumin reservoir

Once absorbed, tirzepatide is more than 99% bound to plasma **albumin**, reversibly, through its C20 fatty-diacid side chain (Chapter 2). The volume of distribution is small, reflecting that the drug is largely confined to plasma and extracellular fluid rather than widely distributed into tissues. This albumin binding is the central pharmacokinetic fact of the molecule: simultaneously the distribution reservoir, the protection against degradation and renal filtration, and the slow-release depot that sets the long half-life. The small free fraction in equilibrium with the bound pool diffuses to target tissues — including, by the routes of Chapters 6 and 7, the relevant brain regions — to engage the GIP and GLP-1 receptors.

## 8.3 Half-life, steady state and the missed-dose rule

The elimination half-life of tirzepatide is approximately **five days**, slightly shorter than semaglutide's but still amply long for once-weekly dosing. Three clinically important consequences follow. First, steady state is reached after roughly four to five half-lives, so plasma concentrations continue to rise for about **four weeks** after starting or after each dose increase — which is why the effect and tolerability at a given dose cannot be fully judged for several weeks, and why titration is deliberately gradual (Chapter 10). Second, the multi-day half-life buffers a **missed dose**: a missed weekly dose can be taken within about four days of the scheduled day, after which the patient resumes the regular schedule. Third, the same slow decline means effects — including adverse effects, and any action relevant to pregnancy — **persist for several weeks after the last dose**; there is no rapid washout, so discontinuation before a planned pregnancy or certain procedures must be timed in advance.

> **PEARL** As with semaglutide, almost every dosing rule restates the half-life. Tirzepatide's ~5-day half-life means slow accumulation over ~4 weeks (gradual titration, delayed judgement of effect) and slow decline (a late dose is tolerated, but the drug lingers for weeks after stopping).

## 8.4 Metabolism — why there are almost no CYP interactions

Tirzepatide is a peptide and is metabolised as one, not by the cytochrome P450 (CYP) system. Elimination proceeds by *proteolytic cleavage* of the peptide backbone by ubiquitous peptidases and by *β-oxidation* of the fatty-acid side chain — the body's ordinary route for fatty-acid breakdown — yielding small peptide fragments and fatty-acid metabolites handled by general catabolic machinery rather than a single organ. The clinical dividend is the same as for semaglutide: because tirzepatide is neither metabolised by nor a meaningful inhibitor or inducer of CYP enzymes, it is essentially free of metabolic drug–drug interactions, requiring no CYP-based screening or dose adjustment of co-administered CYP substrates.

## 8.5 Elimination and special populations

Tirzepatide's metabolites are excreted in urine and faeces, with little intact drug renally cleared because the albumin-bound parent is not filtered (Chapter 2). This underlies a forgiving special-populations profile. **Renal impairment:** clearance does not depend on renal excretion of intact drug, so *no dose adjustment is required* across the range of renal function, including severe impairment, though gastrointestinal tolerability should be watched given

09 SELECTION

# Patient Selection — For Whom, and For Whom Not

The mechanisms of Chapters 3 to 7 dictate the ideal candidate for tirzepatide: someone whose disease is driven by insulin resistance and adiposity — and, given the drug's magnitude of effect and its distinctive sleep-apnoea and heart-failure data, someone with obesity-related complications where that magnitude matters. This chapter sets out who the drug is for, who benefits most, who should not receive it and why, the special situations demanding caution, and the pre-prescription assessment that should precede the first dose. The right patient is the one whose pathophysiology matches the dual mechanism and whose risk of the feared adverse events (Chapter 17) is acceptable.

## 9.1 The established indications

Tirzepatide's indications span glycaemia, weight and — uniquely — a mechanical complication of obesity (Table 9.1). For **type 2 diabetes** (as Mounjaro), it is an adjunct to diet and exercise, drawing on the dual insulin-secretory, glucagon-suppressing, adipose and weight-loss mechanisms of Chapters 3 to 6. For **chronic weight management** (as Zepbound), it is indicated at a BMI of 30 or above, or 27 or above with at least one weight-related comorbidity, on the strength of the SURMOUNT trials and the largest weight loss of the licensed incretins (Chapter 12). Most distinctively, tirzepatide is the **first drug ever approved for obstructive sleep apnoea** (moderate-to-severe OSA with obesity, on the basis of SURMOUNT-OSA, Chapter 15) — a landmark demonstrating that removing the fat substrate of a mechanical airway disease can treat it pharmacologically. And in **heart failure with preserved ejection fraction and obesity**, the SUMMIT trial (Chapter 15) supports an emerging indication. These indications are the clinical harvest of the integrated phenotype (Chapter 6) amplified by the dual mechanism's greater potency.

**Table 9.1. Approved/established indications (tirzepatide).**

| Indication | Threshold / population | Evidence |
|---|---|---|
| Type 2 diabetes | Adjunct to diet/exercise (Mounjaro) | SURPASS[4,9] |
| Chronic weight management | BMI ≥ 30, or ≥ 27 with ≥ 1 comorbidity (Zepbound) | SURMOUNT-1/2[5,7] |
| Obstructive sleep apnoea | Moderate–severe OSA with obesity (first drug approved, Dec 2024) | SURMOUNT-OSA[13] |
| Heart failure (HFpEF + obesity) | Emerging | SUMMIT[12] |

## 9.2 Who benefits most

The patients who benefit most are those whose clinical picture is dominated by the mechanisms the drug addresses, with a particular tilt toward those needing the *greatest* effect. The patient with **type 2 diabetes and obesity** gains glycaemic control and substantial weight loss together. The patient for whom **the magnitude of weight loss is itself the goal**

— severe obesity, or obesity with multiple weight-driven complications — is well served by tirzepatide's class-leading efficacy (Chapter 12). The patient with **obesity-related obstructive sleep apnoea** has, in tirzepatide, the first pharmacological option for that condition. The patient with **obesity and HFpEF** draws on the SUMMIT evidence. And, as for the class, the patient with **a craving-driven eating pattern** (Chapter 7) benefits from reduced food reward. Benefit is largest where insulin resistance and adiposity dominate; where hyperglycaemia is driven purely by advanced beta-cell failure, the insulin-secretory contribution is smaller (Chapter 3) and expectations should be adjusted.

## 9.3 Contraindications and major cautions

The situations to avoid or approach with caution (Table 9.2) follow from mechanism and adverse-effect profile (Chapter 17). The single formal class **contraindication** is a personal or family history of **medullary thyroid carcinoma** (MTC) or **MEN-2**, derived from rodent C-cell tumour findings; although human population data have been reassuring (Chapter 17), the boxed contraindication remains and should be respected. Among **cautions**: prior **pancreatitis** warrants a weighing of risk; **severe gastroparesis** or gut-motility disorder is a particular concern for tirzepatide, which carries the *highest vomiting signal of the class* (Chapter 17), so compounding delayed emptying is best avoided; and a prior serious hypersensitivity reaction contraindicates rechallenge. As with semaglutide, attention to diabetic retinopathy and the rate of glucose lowering is prudent in patients with diabetes (Chapter 16). Each caution is a prompt to assess individual baseline risk, not an automatic bar.

**Table 9.2. For whom *not* — contraindications and major cautions.**

| Situation | Reason |
|---|---|
| Personal/family history of medullary thyroid carcinoma or MEN-2 | Rodent C-cell tumour signal; class boxed contraindication |
| Pregnancy / planning pregnancy / breastfeeding | Inadequate safety data; stop ≥ 2 months before planned conception |
| Prior pancreatitis | Use caution; weigh risk (Chapter 17) |
| Severe gastroparesis / gut-motility disorder | Delayed emptying compounded; highest vomiting signal of the class |
| Type 1 diabetes / DKA | Not a substitute for insulin; not indicated |
| Oral contraceptive users | Counsel on possible reduced efficacy around initiation/escalation; consider a barrier method or non-oral contraception |

## 9.4 Type 1 diabetes and insulin-deficient states

As with all incretin agonists, and following directly from the mechanism of Chapter 3, tirzepatide is **not a substitute for insulin**. Its glucose-lowering depends on amplifying insulin secretion from functioning beta cells and suppressing glucagon; in **type 1 diabetes**, where beta cells are largely absent, that mechanism cannot operate, and using tirzepatide in place of insulin risks **diabetic ketoacidosis**. It is not indicated as insulin replacement and

10 DOSING I

# Subcutaneous Titration

As with all incretin agonists, the governing principle of tirzepatide dosing is that **tolerability — not efficacy — sets the pace**. The schedule is deliberately slow because the gastrointestinal signal is dose-related and concentrated during escalation, and because tirzepatide carries the highest vomiting signal of the class (Chapter 17), unhurried titration matters even more here than for a selective GLP-1 agonist. This chapter sets out the escalation schedule, the administration technique, how to manage the titration phase, missed doses and switching, and the co-medication adjustments — including the tirzepatide-specific oral-contraceptive re-counselling — that must accompany initiation and each dose increase.

## 10.1 Why titration is mandatory

The need to start low and go slow rests on the same mechanistic facts as for any incretin agonist (and is developed fully in the semaglutide volume). The dose-limiting effects — nausea, vomiting and the other gastrointestinal symptoms of Chapter 17 — arise from the same actions that produce benefit: engagement of the hindbrain area postrema (the accessible emetic and satiety sensor, Chapter 6) and slowing of gastric emptying. These effects are dose-related and show **tachyphylaxis**, attenuating as the patient adapts; gradual escalation exploits that adaptation, each tolerated step enabling the next. The pharmacokinetics (Chapter 8) reinforce the schedule: with a ~5-day half-life, each dose takes about four weeks to approach steady state, so a roughly four-week interval per step both allows adaptation and lets each dose near its full effect before the next increase. One nuance is specific to the dual agonist: the GIP arm may itself temper nausea (Chapter 6), which is part of why tirzepatide can ultimately be escalated to very high, highly effective doses — but this does not remove the need for slow titration, it makes the slow climb achievable.

## 10.2 The standard schedule

Tirzepatide is given once weekly by subcutaneous injection, escalating in approximately four-week steps: **2.5 → 5 → 7.5 → 10 → 12.5 → 15 mg**. The **2.5 mg dose is an initiation (tolerability) dose only** — not a maintenance dose; its purpose is to begin the adaptation, not to treat. Maintenance is typically 5, 10 or 15 mg depending on response and tolerability. As with semaglutide, the cardinal rule is to *escalate only when the current dose is adequately tolerated*: if symptoms are marked at a step, hold (or step down) for additional weeks and re-attempt the increase once symptoms settle, rather than forcing progression. The four-week interval is a minimum, not a mandate.

**PEARL** Many patients reach target weight or glycaemic goals below 15 mg. There is no obligation to climb to the maximum dose; the lowest dose that achieves the goal with good tolerability is the right maintenance dose.

## 10.3 Administration technique

Tirzepatide is injected once weekly into the abdomen, thigh or upper arm, with site rotation to reduce local reactions; the same weekday should be kept for consistency, though it may be changed when needed provided at least about three days separate consecutive doses. The device requires appropriate cold-chain storage before first use, with limited time at room temperature permitted thereafter per the product instructions; the peptide is sensitive to heat and freezing, and a device that has been frozen or overheated should not be used (Chapter 2). There is no oral formulation and therefore no empty-stomach ritual to teach (Chapter 8). Patient education at initiation should cover injection technique, the expected pattern of early gastrointestinal symptoms, and the contraception counselling of Section 10.6.

## 10.4 Managing the titration phase

The titration phase is where adherence is most fragile, and for tirzepatide — with the highest vomiting signal of the class — proactive management is especially valuable. The pace should be individualised: hold or step down rather than forcing escalation through significant symptoms, with reassurance that the symptoms of a given step typically settle over its weeks. Simple dietary measures (smaller, lower-fat meals; eating slowly; stopping at fullness) and, where appropriate, short-term antiemetic support reduce nausea (Chapters 17 and 18). Adequate hydration is important if vomiting or diarrhoea occurs, both for comfort and to protect renal function (Chapter 17). Setting expectations in advance — that early nausea is common, dose-related and self-limiting — markedly improves the odds of reaching an effective maintenance dose.

## 10.5 Missed doses and switching

The forgiving pharmacokinetics (Chapter 8) make missed doses manageable: a missed weekly dose can be taken within about four days of the scheduled day, after which the patient resumes the regular schedule; beyond that window the missed dose is skipped and the next is taken on the regular day. After a prolonged interruption (several weeks off therapy), tolerability adaptation may be lost, and it is prudent to re-initiate at a lower step and re-titrate rather than resuming at the previous maintenance dose. Switching between tirzepatide and a GLP-1 receptor agonist (such as semaglutide) is *not* dose-equivalent: the new agent should be started on its own initiation step and titrated on its own schedule, since tolerance of a high dose of one does not license starting at a high dose of the other.

## 10.6 Co-medication adjustment — glucose-lowering drugs and contraception

Titration must be accompanied by two co-medication adjustments. First, as for any incretin agonist and following from the mechanism of Chapter 3, a background **sulfonylurea or insulin** — which lowers glucose independently of the prevailing level — should be *reduced pre-emptively* at tirzepatide initiation and reviewed at each escalation, to avoid hypoglycaemia; metformin and most other agents do not require this. Second, and specific to tirzepatide, patients using **oral hormonal contraceptives** should be counselled at *initiation and at each dose increase* that the delayed gastric emptying can reduce contraceptive efficacy, and advised to use a non-oral method or add barrier protection for about four weeks after

11 DOSING II

# Minimum Effective Dose, Plateaus, and the Current Landscape

Chapter 10 set out how to escalate tirzepatide; this chapter addresses the questions that follow once a patient is established: what is the right dose for *this* patient and goal, how low can a maintenance dose be, why does weight loss stall, and how to use the dose to break a plateau. It then defines the off-label practices of micro- and macro-dosing and sets out the mid-2026 landscape. The unifying theme is *reason dosing* — choosing the dose by goal, response and tolerability rather than reflexively maximising or minimising it.

## 11.1 The minimum effective dose principle

The most useful dosing principle is that the goal is the **minimum effective dose** — the lowest dose achieving the patient's objective with acceptable tolerability — not the maximum the label permits. The titration schedule of Chapter 10 (2.5 → 5 → 7.5 → 10 → 12.5 → 15 mg) climbs toward a ceiling, but the ceiling is a *limit*, not a *target*. The dose–response relationship is real but its benefit-to-harm ratio is not linear: higher doses produce greater average weight loss and glycaemic effect (SURMOUNT-1 showed roughly 15%, 19.5% and 20.9% at 5, 10 and 15 mg respectively), but also more gastrointestinal toxicity, with diminishing incremental benefit and undiminished incremental burden at the top of the range. Many patients reach their goal at 5 or 10 mg without needing 15 mg, and for them there is no clinical reason — and a real tolerability cost — to climb further. The corollary, *reason dosing*, is that every dose decision should have an articulable reason: escalate because the goal is unmet and the dose is tolerated; hold because the goal is met or tolerability is marginal; or hold at the lowest dose that maintains the result.

For *maintenance*, once a patient reaches and stabilises at their goal, the aim becomes the lowest dose that *holds* the result with good tolerability. Honesty is required here: the trials studied fixed target doses, and evidence for deliberately down-titrating after success is limited, with a recognised risk that reducing the dose allows partial regain as the defended set-point reasserts itself (Chapters 6 and 18). The defensible practice is to individualise — the lowest dose that demonstrably maintains the goal for that patient, monitored for regain — not to assume everyone must remain at 15 mg or that everyone can be safely reduced.

**PEARL — reason dosing**
15 mg is a ceiling, not a target. Ask of every dose: *what is the clinical reason this patient is on this dose?* Goal met and tolerated → no reason to climb. Goal unmet, dose tolerated and below 15 mg → a reason to climb. The plateau question is this principle applied to a stall.

## 11.2 The weight-loss plateau — why it happens

Almost every patient eventually reaches a **weight-loss plateau** — weight stabilises and further loss stops — and understanding why is essential to responding correctly. In most cases the plateau is a predictable physiological equilibrium, not a failure. As weight falls, two homeostatic responses (Chapter 6) progressively oppose further loss: total energy expenditure

declines (a smaller body burns fewer calories, plus **adaptive thermogenesis**, a disproportionate fall in expenditure), and the defended set-point partially reasserts itself so appetite recovers somewhat from its suppressed state. As intake and expenditure re-equilibrate at the new, lower weight, loss ceases. At a fixed dose this new equilibrium is the expected endpoint, not the drug "stopping working."

Before treating a plateau as a true equilibrium, exclude two *pseudo-plateaus* that call for patience rather than action. First, an apparent stall during titration or in the weeks just after a dose increase may simply mean the current dose has *not yet reached steady state* — each dose takes about four weeks to express its full effect (Chapter 8). Second, an apparent stall in a patient still on a sub-maximal dose is not a true ceiling but a plateau *at the current dose* that the dose–response relationship predicts can be moved by escalation. Distinguishing a true, maximal-dose equilibrium from these pseudo-plateaus is the first step.

## 11.3 Breaking a plateau with dosing

Once a genuine plateau is identified, the dosing approach divides according to whether the patient is already at 15 mg.

**If the patient is below 15 mg and tolerating the current dose**, the primary dosing lever is *escalation to the next step*. Because tirzepatide shows a clear dose–response for weight loss, a patient who has plateaued at, say, 5 or 10 mg and is tolerating it can reasonably be advanced to the next dose, with a renewed period of loss often following as the higher dose takes effect over the subsequent weeks. This is the cleanest, best-evidenced way to break a plateau: the stall signals that the current dose has delivered what it can, and dose-response headroom remains. Escalation still follows the tolerability rules of Chapter 10 — advance only if the current dose is tolerated, allow four to five weeks for the new dose to express its effect, and remember to re-counsel oral-contraceptive users at each increase (Chapter 8).

**If the patient is already at 15 mg**, the dose should *not* be pushed beyond the maximum to chase further loss: there is no trial evidence of benefit above 15 mg, the incremental toxicity is real (Section 11.5), and exceeding the licence is macrodosing. A plateau at 15 mg is the expected pharmacological ceiling, and the levers then are not higher dosing but: *verify adherence* (missed doses); *intensify the non-pharmacological programme* — the dietary, protein and resistance-training measures of Chapter 18, which protect lean mass and can resume fat loss; *reassess the goal*, since a plateau at the maximum dose may simply be the patient's new, acceptable maintenance weight; and recognise that, because tirzepatide is already the *highest-efficacy licensed incretin*, there is currently no higher-efficacy agent to switch *up* to (the triple agonist retatrutide, which produced greater weight loss in early trials, is not yet approved — Chapter 19). The key discipline is that, at 15 mg, the dosing lever is exhausted, and the response must come from adherence, lifestyle, and reassessment rather than from exceeding the licence.

12 EFFICACY I

# Glycaemic and Weight Efficacy

The dual mechanism of the preceding chapters predicts that tirzepatide should lower glucose and reduce weight more than any selective GLP-1 agonist, and the trial programmes confirm it: tirzepatide is the most effective licensed incretin for both weight and glycaemia. This chapter reviews the weight evidence (the SURMOUNT programme), the head-to-head comparisons with semaglutide that define its place, the glycaemic evidence (SURPASS), the magnitude in context, and the composition and durability of the effect. The figures place the numbers in visual context.

## 12.1 Weight — the SURMOUNT programme

The **SURMOUNT** programme established tirzepatide as the most powerful licensed weight-loss therapy. In the pivotal **SURMOUNT-1** trial, adults with obesity without diabetes lost approximately **15%, 19.5% and 20.9%** of body weight at the 5, 10 and 15 mg doses respectively over 72 weeks — a clear dose–response (Figure 12.1) and, at the top dose, a mean loss exceeding 20%, territory previously the preserve of bariatric surgery. **SURMOUNT-2** studied patients with type 2 diabetes, in whom loss was characteristically somewhat smaller (around 13–15%) but still exceeding what selective agonists achieve in that population. **SURMOUNT-3** used an intensive lifestyle lead-in and **SURMOUNT-4** was a withdrawal trial (Section 12.5). Across the programme, large proportions of patients reached the 15%, 20% and even 25% thresholds at which obesity-related comorbidities substantially improve — the clinical correlate of the dose-dependent curve shown below.

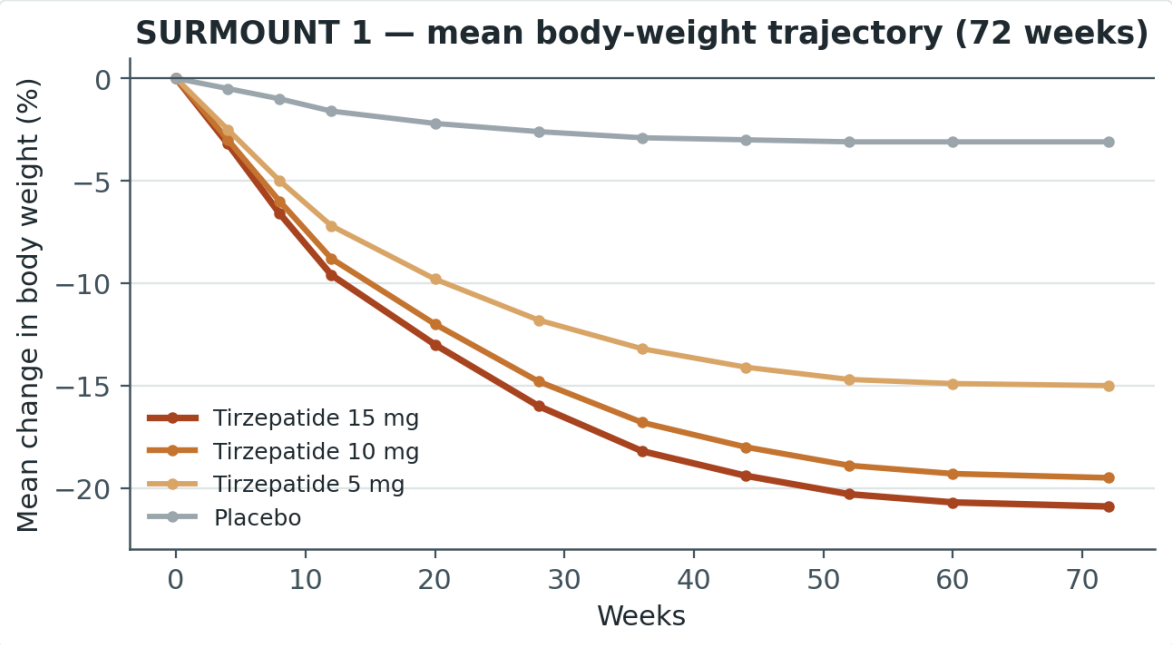



**Figure 12.1.** SURMOUNT-1: dose-dependent weight loss with tirzepatide 5/10/15 mg versus placebo over 72 weeks.

## 12.2 Head-to-head with semaglutide

The comparisons that most clearly define tirzepatide's place are the direct head-to-head trials against semaglutide, the most potent selective GLP-1 agonist. In **SURMOUNT-5** (obesity), tirzepatide produced greater weight loss — around 20% versus about 14% for semaglutide 2.4 mg — and a higher proportion of patients reaching every weight-loss threshold (Figure 12.2). In **SURPASS-2** (type 2 diabetes), tirzepatide outperformed semaglutide 1.0 mg on both HbA1c and weight, across its dose range (Figure 12.3). These results are the clinical expression of the dual-mechanism advantage developed in Chapters 1 to 6 — the added GIP adipose arm and the imbalanced, biased pharmacology that allows escalation to high efficacy with comparable tolerability. They establish tirzepatide as the efficacy leader among licensed incretins for both endpoints.

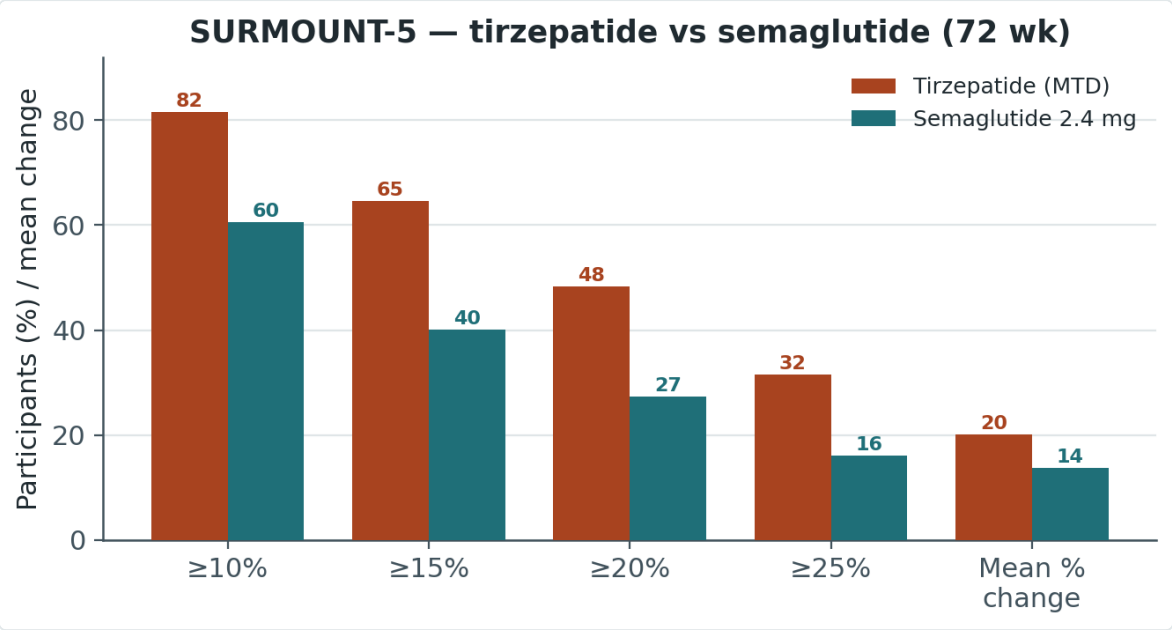



**Figure 12.2.** SURMOUNT-5: proportions reaching weight-loss thresholds, tirzepatide versus semaglutide 2.4 mg.

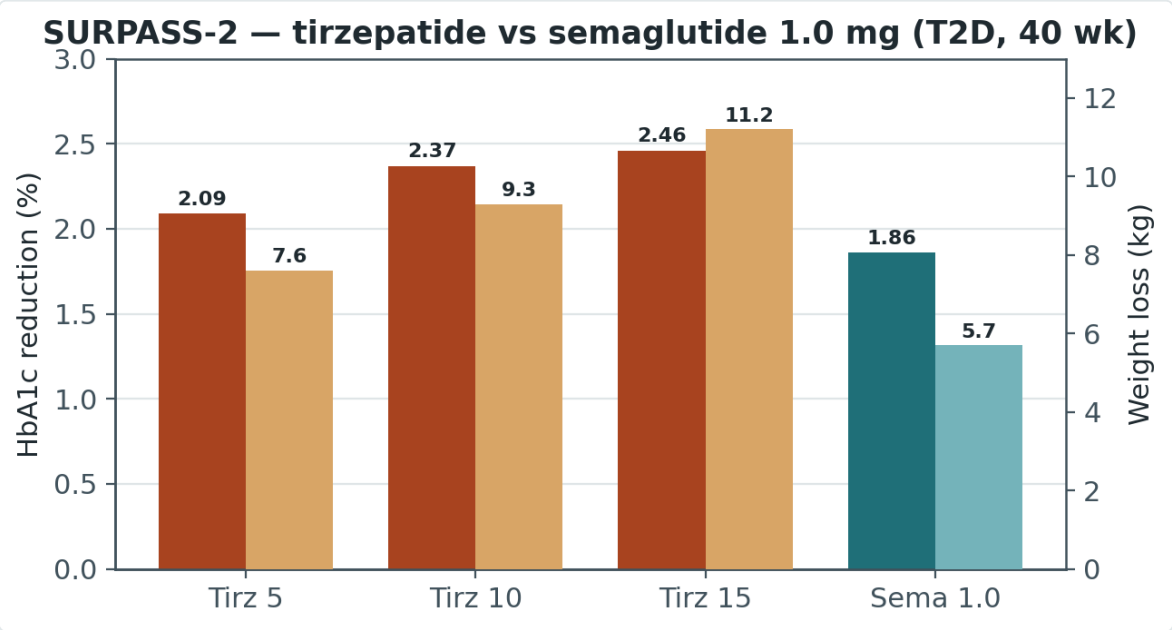



**Figure 12.3.** SURPASS-2: HbA1c and weight reductions, tirzepatide versus semaglutide 1.0 mg in type 2 diabetes.

13 OUTCOMES I

# Cardiovascular Outcomes and Their Mechanism

Glycaemic and weight efficacy (Chapter 12) matter because of what they prevent. For cardiovascular outcomes, tirzepatide's evidence is genuinely strong but, in an important respect, *less mature* than semaglutide's — a distinction a prescriber should understand clearly rather than assume the dual agonist's superior metabolic effects automatically translate into superior proven outcomes. This chapter reviews the cardiovascular trial evidence, states honestly what has and has not yet been demonstrated, and develops the mechanisms — including the distinctive contribution of the GIP adipose arm — that are expected to underlie benefit.

## 13.1 The evidence

**SURPASS-CVOT** tested tirzepatide against the active comparator *dulaglutide* — an agent already proven to reduce cardiovascular events — in type 2 diabetes with cardiovascular disease. It met non-inferiority for MACE with a hazard ratio of approximately **0.92**, confirming cardiovascular safety and a benefit at least matching an established GLP-1 agonist, while weight, HbA1c and blood pressure improved more with tirzepatide. Separately, the **SUMMIT** trial (developed in Chapter 15) showed benefit in heart failure with preserved ejection fraction and obesity — a cardiovascular-adjacent outcome of real importance given how common and undertreated that condition is. A dedicated obesity cardiovascular-outcomes trial, **SURMOUNT-MMO**, is ongoing and will test event reduction in obesity without diabetes, as SELECT did for semaglutide. Figure 13.1 places SURPASS-CVOT alongside the semaglutide trials for context.

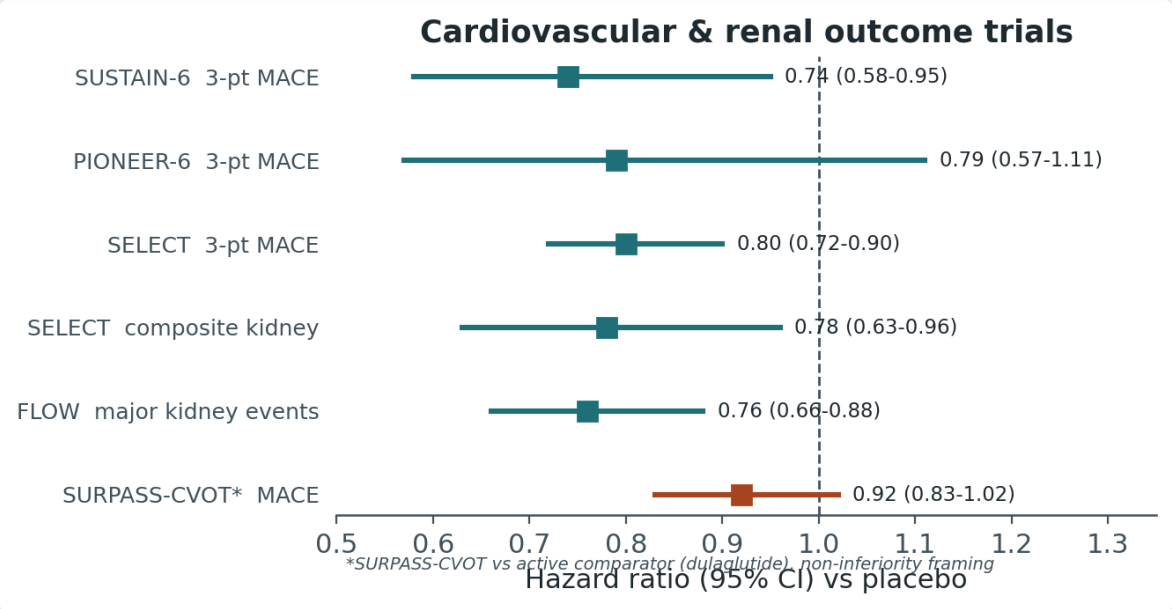



**Figure 13.1.** Incretin cardiovascular/renal outcome trials. SURPASS-CVOT (tirzepatide) is versus active comparator (dulaglutide); semaglutide trials shown for context.

## 13.2 What has, and has not, been shown — an honest reading

Precision about the evidence matters here. SURPASS-CVOT was an *active-comparator* trial designed to show non-inferiority to dulaglutide, not a placebo-controlled superiority trial; it therefore establishes that tirzepatide is cardiovascular-safe and at least as protective as a proven GLP-1 agonist, but it does not, by itself, quantify a placebo-referenced event reduction of the kind SELECT provided for semaglutide in obesity. The dedicated obesity outcomes trial (SURMOUNT-MMO) that would furnish that placebo-referenced, weight-independent evidence has not yet reported. The reasonable current position is therefore: tirzepatide is cardiovascular-safe, at least matches an established GLP-1 agonist on hard events while improving the metabolic profile more, and shows benefit in HFpEF (SUMMIT) — but the mature, placebo-controlled, obesity cardiovascular-outcome evidence that semaglutide already holds (SELECT) is, for tirzepatide, still awaited. This is not a mark against the drug so much as a matter of evidentiary timing, and it is precisely the kind of distinction that should inform agent choice when the goal is proven cardiovascular protection (Chapter 9).

## 13.3 Why — mechanisms of cardiovascular benefit

The mechanisms expected to underlie tirzepatide's cardiovascular benefit overlap substantially with those established for GLP-1 agonism (Table 13.1), with an added dimension from the GIP arm. As with semaglutide, benefit is expected to be only *partly* weight-mediated, operating also through direct anti-inflammatory, endothelial and haemodynamic actions — though the definitive weight-independent obesity data await SURMOUNT-MMO.

**Table 13.1. Mechanisms of cardiovascular benefit.**

| Pathway | Effect |
|---|---|
| Anti-inflammatory | Reduced systemic inflammation (hsCRP) and, by class, attenuated atherosclerotic gene programmes[17] |
| Endothelial / vascular | Improved endothelial function and vascular reactivity |
| Haemodynamic / metabolic | Lower blood pressure, improved lipids, reduced visceral and epicardial fat |
| GIP adipose arm | Improved lipid handling and insulin sensitivity (Chapter 5) reduce atherogenic dyslipidaemia |
| Glycaemic | Large HbA1c reduction and high rates of normoglycaemia |

The distinctive strand is the **GIP adipose arm** (Chapter 5). By improving how adipose tissue buffers dietary lipid and by enhancing insulin sensitivity, the GIP component reduces the atherogenic dyslipidaemia and ectopic-fat burden that drive cardiovascular risk — potentially adding a metabolic route to cardiovascular benefit beyond that of a selective GLP-1 agonist. Whether this translates into greater *hard-outcome* cardiovascular benefit than semaglutide is exactly the question SURMOUNT-MMO is designed to answer; mechanistically it is plausible, but it should not be asserted as established until the outcome data exist. The reduction of visceral and epicardial fat, the improvement in blood pressure and lipids, and the large glycaemic effect complete a coherent picture in which the greater metabolic effect of the dual agonist is expected to confer at least comparable, and possibly greater, cardiovascular protection — pending confirmation.

14 OUTCOMES II

# Renal Outcomes and Their Mechanism — SURPASS-4

Diabetic kidney disease is a leading cause of kidney failure and a powerful multiplier of cardiovascular risk. The incretin class as a whole slows its progression, and tirzepatide shows a concordant renal signal — but, as with the cardiovascular data of Chapter 13, the renal evidence for tirzepatide is at present *less mature* than semaglutide's, resting on a pre-specified analysis rather than a dedicated kidney-outcome trial. This chapter reviews that evidence honestly and develops the mechanisms of renal protection, including the GIP-related contribution.

## 14.1 The evidence

A pre-specified renal analysis of **SURPASS-4** — which studied patients with type 2 diabetes at high cardiovascular risk — found that tirzepatide **slowed the decline in eGFR**, **reduced new-onset macroalbuminuria**, and lowered a composite kidney endpoint compared with insulin glargine. These are meaningful, disease-relevant renal signals, and they mirror the established renal benefit of GLP-1 agonists. It is important, however, to be precise about the evidentiary status: SURPASS-4's renal findings come from a *secondary, pre-specified analysis versus an active comparator*, not from a dedicated, adequately powered, placebo-referenced kidney-outcome trial of the kind **FLOW** provided for semaglutide. A dedicated tirzepatide kidney-outcome trial is the logical next step and would be needed to establish a placebo-referenced hazard ratio for hard renal endpoints. The reasonable current position is that tirzepatide shows a renal-protective signal consistent with its class and its metabolic effects, with the definitive dedicated-trial evidence still to come.

## 14.2 Why — mechanisms of renal protection

The mechanisms expected to underlie tirzepatide's renal benefit parallel those of the class (Table 14.1), with the metabolic advantages of the dual agonist reinforcing several of them.

**Table 14.1. Mechanisms of renal benefit.**

| Pathway | Effect |
|---|---|
| Reduced glomerular hyperfiltration | Lower intraglomerular pressure → less mechanical injury and falling albuminuria |
| Natriuresis | Reduced proximal-tubule sodium reabsorption; lower volume and blood pressure |
| Anti-inflammatory / anti-fibrotic | Reduced renal inflammatory and fibrotic signalling |
| Systemic risk-factor control | Lower BP, glucose and weight reduce drivers of progression |

The **haemodynamic** arm is probably central. In diabetes and obesity the glomerulus is subjected to **hyperfiltration** — an elevated intraglomerular pressure that mechanically injures podocytes and the capillary wall, driving albuminuria and scarring. Incretin-receptor

signalling reduces this hyperfiltration, partly through a *natriuretic* action: reduced proximal-tubule sodium reabsorption increases distal sodium delivery, which via tubuloglomerular feedback constricts the afferent arteriole and lowers intraglomerular pressure (mechanistically analogous to, and complementary with, SGLT2 inhibitors). Direct **anti-inflammatory and anti-fibrotic** renal actions and broad **systemic risk-factor control** (lower glucose, blood pressure, weight and improved lipids) add further protection. Tirzepatide's distinctive contributions are its *larger weight loss* — which more powerfully relieves obesity-related glomerular hyperfiltration and glomerulopathy — and its *GIP-driven improvement in insulin sensitivity and lipid handling* (Chapter 5), which reduces the metabolic drivers of renal injury. Whether these translate into greater hard renal-outcome benefit than a selective GLP-1 agonist is, again, a question only a dedicated outcome trial can answer.

## 14.3 Clinical implications

For a patient with type 2 diabetes and chronic kidney disease, tirzepatide's renal signal and its powerful metabolic effects support its use as part of a comprehensive strategy, alongside RAS blockade, an SGLT2 inhibitor and, where appropriate, finerenone. The pharmacokinetics support this: because tirzepatide is cleared by metabolism rather than renal excretion of intact drug, *no dose adjustment is required* across the range of renal function (Chapter 8). Two cautions apply, and are if anything more important for tirzepatide given its higher gastrointestinal signal (Chapter 17): gastrointestinal effects causing vomiting or diarrhoea can precipitate dehydration and a pre-renal fall in kidney function, so hydration and careful titration matter especially in patients with existing renal impairment; and the doses of other glucose-lowering agents (insulin, sulfonylureas) should be reviewed as glucose falls and as renal impairment alters their clearance. For the specific goal of *proven* renal-outcome protection, semaglutide currently holds the more mature evidence (FLOW); tirzepatide offers a concordant signal with superior metabolic effect, pending its own dedicated trial — a distinction that, as in Chapter 13, should inform agent choice when renal protection is the primary aim.

### ◆ KEY POINTS — CHAPTER 14

- **SURPASS-4 renal analysis:** slowed eGFR decline, less new macroalbuminuria, lower composite kidney endpoint vs glargine — a concordant renal signal.
- **Honest reading:** this is a pre-specified, active-comparator secondary analysis, not a dedicated placebo-referenced kidney-outcome trial (as FLOW was for semaglutide); a dedicated trial is still needed.
- Mechanisms parallel the class — reduced hyperfiltration via natriuresis/tubuloglomerular feedback, anti-inflammatory/anti-fibrotic effects, systemic risk-factor control — reinforced by tirzepatide's larger weight loss and the GIP metabolic arm.
- No renal dose adjustment (Chapter 8); watch dehydration-related pre-renal decline (higher GI signal) and review insulin/sulfonylurea dosing.
- For *proven* renal outcomes, semaglutide is currently better evidenced; tirzepatide offers a concordant signal with greater metabolic effect, pending a dedicated trial.

15 OUTCOMES III

# Liver, Heart Failure and Sleep Apnoea

This chapter is one of tirzepatide's strengths, because its greater magnitude of weight loss translates into greater organ-specific benefit across the diseases of adiposity — and because it holds a landmark that no other drug does: the first approval for obstructive sleep apnoea. Metabolic liver disease, obesity-related HFpEF and OSA are three organ diseases of ectopic and excess fat, and tirzepatide's effect on each follows from removing the fat and inflammation that drive them. This chapter reviews the evidence in each, with particular attention to the sleep-apnoea data that make tirzepatide unique.

## 15.1 Liver — MASLD / MASH

MASLD and its progressive form MASH are the hepatic face of ectopic fat: lipid in the hepatocyte (Chapter 5) drives inflammation and, over time, fibrosis and cirrhosis, and MASH has become a leading cause of chronic liver disease with little effective pharmacotherapy until recently. The **SYNERGY-NASH** trial tested tirzepatide in biopsy-confirmed MASH and showed **dose-dependent resolution of steatohepatitis** — on the order of **52%, 63% and 73%** at the 5, 10 and 15 mg doses respectively, versus a low placebo rate — among the highest resolution rates reported for any pharmacotherapy, with improvement in fibrosis markers and no worsening of fibrosis. The mechanism is the ectopic-fat depletion of Chapter 5 applied to the liver, reinforced by the drug's systemic anti-inflammatory and insulin-sensitising actions and, distinctively, by the GIP adipose arm that improves lipid handling and reduces the overflow of fat to the liver. As with semaglutide, resolution of the *steatohepatitis* (inflammatory activity) is achieved more readily than regression of established *fibrosis*, but the high resolution rates and the dose–response are a major advance in a disease that badly needed one.

## 15.2 Heart failure with preserved ejection fraction

A large and increasingly recognised subset of heart failure with preserved ejection fraction (HFpEF) is fundamentally *obesity-related*, driven by the volume expansion, systemic inflammation, epicardial fat and haemodynamic burden of obesity. The **SUMMIT** trial tested tirzepatide in obesity-related HFpEF and showed benefit on a composite of cardiovascular death and worsening heart-failure events, together with improvements in **heart-failure symptoms and physical limitation** (KCCQ), **exercise capacity** and inflammatory markers. This is a step beyond the symptom-and-function improvements of the semaglutide HFpEF trial, in that SUMMIT included harder heart-failure event outcomes — reinforcing the reframing of a subset of HFpEF as a treatable consequence of obesity. The mechanism is coherent with the rest of the monograph: reducing weight, volume, visceral and epicardial fat and systemic inflammation relieves the specific pathophysiology obesity imposes on the failing-but-preserved ventricle, and tirzepatide's larger weight loss plausibly drives a correspondingly larger effect.

## 15.3 Obstructive sleep apnoea — the landmark

Obstructive sleep apnoea (OSA) is a mechanical disease of adiposity: fat deposition around the pharyngeal airway, together with the effects of obesity on lung volumes and respiratory mechanics, narrows and destabilises the upper airway during sleep, producing the repetitive collapse, oxygen desaturation and arousals that define the condition and drive its cardiovascular and neurocognitive consequences. The **SURMOUNT-OSA** trial tested tirzepatide in patients with moderate-to-severe OSA and obesity and demonstrated a **large reduction in the apnoea–hypopnoea index** (the measure of respiratory events per hour of sleep), with many patients improving by a full severity category and some reaching resolution thresholds, alongside improvements in oxygenation, blood pressure and inflammatory markers. On the strength of these results, in December 2024 tirzepatide became the **first drug ever approved for the treatment of obstructive sleep apnoea** — a genuine landmark, because OSA had previously been treated only mechanically (continuous positive airway pressure) or surgically, with no effective pharmacological option.

The mechanism is the most direct illustration in the monograph of "removing the substrate treats the disease." By depleting the peri-pharyngeal and visceral fat and reducing the mechanical and inflammatory burden on the airway, tirzepatide addresses the physical cause of the airway collapse rather than merely splinting it open as CPAP does. This makes it, for suitable patients, a disease-modifying rather than purely supportive therapy — and for the substantial number of patients who cannot tolerate CPAP, it offers an effective alternative or adjunct for the first time. The magnitude of tirzepatide's weight loss (Chapter 12) is what made this achievable, which is why the OSA indication belongs to the dual agonist rather than to a selective GLP-1 agonist. Beyond these three systems, tirzepatide's benefits extend, like the class, to conditions such as PCOS and the broader metabolic complications of obesity, following the same unifying principle of Chapter 6.

### ◆ KEY POINTS — CHAPTER 15

- **Liver (MASH):** SYNERGY-NASH showed dose-dependent steatohepatitis resolution ~52/63/73% at 5/10/15 mg — among the highest reported — via hepatic ectopic-fat depletion reinforced by the GIP adipose arm; inflammation resolves faster than fibrosis.
- **HFpEF:** SUMMIT improved a composite of CV death/worsening HF events plus symptoms and exercise capacity in obesity-related HFpEF — harder outcomes than the semaglutide HFpEF trial.
- **OSA — landmark:** SURMOUNT-OSA achieved large apnoea–hypopnoea index reductions, making tirzepatide the **first drug ever approved for obstructive sleep apnoea** (Dec 2024).
- The OSA mechanism — removing peri-pharyngeal/visceral fat to treat the physical cause of airway collapse — is the clearest example of "removing the substrate treats the disease," and offers an option for CPAP-intolerant patients; the large weight loss is why this indication is tirzepatide's.

16 SAFETY — THE EYE

# Diabetic Retinopathy and the Optic-Neuropathy Question

Two distinct ocular issues are frequently conflated — one a treatable consequence of therapeutic *success*, the other a rare ischaemic event of uncertain causal link. For tirzepatide, both are best understood as questions about *rapid glucose lowering and the incretin class* rather than about a unique property of the molecule, and the honest position is that much of the specific evidence (especially for the optic-neuropathy signal) derives from semaglutide, with tirzepatide-specific data thinner. This chapter develops the mechanism and the clinical approach for each.

## 16.1 Early worsening of diabetic retinopathy

Trials of incretin therapy in diabetes have reported **early worsening** of retinopathy concentrated in patients with **pre-existing retinopathy, long-standing poorly controlled diabetes, and the largest, fastest HbA1c falls**. This is not a drug-specific toxin but a general consequence of *rapid glycaemic improvement*: the same phenomenon was described decades ago with intensive insulin therapy, where tight control transiently worsened retinopathy before conferring substantial long-term benefit. The relevance to tirzepatide is direct and, if anything, heightened: because tirzepatide produces the *largest* HbA1c reductions of the incretin class (Chapter 12), a patient with poorly controlled diabetes and established retinopathy started on tirzepatide may experience a particularly large and rapid fall in glucose — exactly the setting in which early worsening occurs. This makes the preventive approach below especially pertinent for the dual agonist.

## 16.2 The mechanism of early worsening

The retina of a chronically hyperglycaemic patient has adapted to its abnormal environment; sudden correction perturbs several systems at once. Retinal blood flow shifts as glucose falls, and the autoregulation of the already-compromised retinal vasculature copes poorly with the change. Growth-factor signalling is also perturbed — rapid metabolic change can transiently alter insulin-like growth factor and vascular endothelial growth factor (VEGF), the driver of the pathological neovascularisation of proliferative retinopathy — in ways that can briefly accelerate existing disease. Eyes with *established* retinopathy have the vascular fragility and neovascular tendency for these perturbations to act upon and can worsen in the first months of rapid improvement; eyes without significant baseline retinopathy have little for the phenomenon to act on and are at low risk. This explains the risk profile precisely — the danger is concentrated in the poorly controlled patient with existing retinopathy undergoing a large, fast fall in HbA1c — and, given tirzepatide's potency, reinforces the value of a graded approach in such patients.

**PEARL** Because tirzepatide lowers HbA1c the most of the class, the early-worsening risk is most relevant here. In a patient with diabetes, poor control and known retinopathy, arrange a baseline dilated fundus exam, aim for more gradual glucose lowering, and co-manage with ophthalmology. The at-risk patient is identifiable in advance.

## 16.3 The optic-neuropathy (NAION) question

Non-arteritic anterior ischaemic optic neuropathy (NAION) presents as sudden, usually painless, monocular vision loss from impaired perfusion of the optic-nerve head; the loss is generally irreversible. It is an entirely different entity from retinopathy — an ischaemic event of the optic nerve, not a microvascular retinal disease — and its possible association with incretin therapy emerged recently and remains genuinely uncertain. The specific evidence to date centres on *semaglutide*: following a 2024 single-centre signal, several cohorts and a meta-analysis found roughly a **two-fold** association (pooled hazard ratio around 2.6, significant mainly after about two years), while other matched analyses found **no** excess, and in June 2025 the European Medicines Agency concluded that NAION is a **very rare** side effect of semaglutide (up to about 1 in 10,000). For *tirzepatide specifically*, dedicated data are much more limited, and whether the signal is a class effect, a semaglutide-specific finding, or largely an artefact of confounding is not established.

The interpretive difficulty is a textbook case of **confounding by indication**: the risk factors for NAION — older age, male sex, long diabetes duration, high HbA1c, existing retinopathy, obesity, and a small "crowded" optic disc — overlap almost entirely with the population that takes these drugs, so an association appears in observational data whether or not the drug is causal. The honest position for tirzepatide is therefore one of proportionate, class-informed caution: the absolute risk, if any, is very low; the tirzepatide-specific evidence is thin; and the appropriate response is awareness rather than alarm, pending better data.

## 16.4 A rational clinical approach to the eye

For **early retinopathy worsening**, the approach is preventive and targets the identifiable at-risk patient — and is especially pertinent given tirzepatide's potency: in a person with poorly controlled diabetes and known or likely retinopathy, arrange a baseline dilated fundus examination around initiation, aim to lower glucose more gradually rather than precipitously, and co-manage with ophthalmology through the first months of improvement. In patients without significant retinopathy or a large, rapid HbA1c fall, routine baseline eye assessment follows standard diabetes-care guidance. For **NAION**, given the low absolute risk, the uncertain causality, and the thin tirzepatide-specific evidence, the proportionate approach is to inform patients — particularly those with overlapping risk factors — to report any sudden change in vision promptly, arrange urgent ophthalmological assessment if it occurs, and discontinue if NAION is confirmed; routine screening specifically for NAION is not warranted. Prevent the predictable problem in the identifiable patient, remain alert to the rare one, and do not let a low-quality, confounded signal deny patients the substantial benefits of therapy.

17 SAFETY IN DEPTH

# Adverse Effects — Mechanisms, Predictors, Prevention and Confounders

Most adverse effects trace to **one of two mechanisms**: the delayed gastric emptying and area-postrema signalling that also drives benefit, or the consequences of **rapid weight loss itself**. A third theme runs throughout: many "side effects" are partly the patient's pre-existing disease, not the drug.

## 17.1 Gastrointestinal effects — the dominant, dose-related class

**Why.** Nausea, vomiting, early satiety, diarrhoea and constipation arise from delayed gastric emptying plus direct area-postrema (emetic-sensor) activation.[33] **Predictors:** faster titration, higher dose, female sex, and the post-escalation period. Tirzepatide carries the **highest vomiting signal** of the class in trial meta-analysis (relative risk markedly above semaglutide), so titration discipline matters even more.[33] **Prevention:** slow titration; small, low-fat, low-volume meals; stop at fullness; hydrate; **hold or step down** rather than push on. The GIP component may modestly temper nausea relative to an equipotent pure GLP-1 agonist, but the class effect dominates.

**Time course and escalation of management.** Nausea peaks in the days after each dose increase and settles over the subsequent weeks as gastric-emptying tachyphylaxis develops (Chapter 6), so the pattern is transient peaks on a declining baseline rather than a steady burden — reassuring to convey to patients. When dietary measures are insufficient, a short course of an antiemetic during the worst days after an increase is reasonable, and persistent or severe symptoms should prompt a dose hold or step-down rather than escalation. Given tirzepatide's **highest-in-class vomiting signal**, this discipline matters more than for any other agent, and the downstream consequences of vomiting and diarrhoea are correspondingly important: dehydration and its renal effects (Section 17.7), and reduced oral-contraceptive absorption (Chapter 8). A minority cannot tolerate the drug at a useful dose despite optimal measures; recognising this early spares a futile, miserable course.

## 17.2 Gallbladder and biliary disease

**Why. Rapid weight loss** supersaturates bile with cholesterol (a risk with weight loss of any cause), and reduced gallbladder contractility promotes bile stasis.[35] **Predictors:** rapid/large loss (tirzepatide produces the largest), higher dose, longer use, female sex, prior stones. **Prevention:** avoid excessively rapid loss; counsel on biliary symptoms; investigate right-upper-quadrant pain.

## 17.3 Pancreatitis — signal, and the confounding problem

A 2025 meta-analysis of 62 randomised trials found a modestly increased relative risk of acute pancreatitis (≈ 1.4) across the class, with **no** association with pancreatic cancer.[34] The qualifying population **already carries pancreatitis risk factors** — gallstones, hypertriglyceridaemia, obesity, alcohol, diabetes — and rapid weight loss is itself a risk factor, so part of the signal is the population, not the molecule. **Action:** avoid where there is a clear

prior history and an alternative exists; stop and investigate severe persistent abdominal pain; do not restart after confirmed drug-related pancreatitis.

## 17.4 Pulmonary aspiration and anaesthesia

**Why.** Delayed gastric emptying leaves **retained gastric contents despite standard fasting**; aspiration has occurred at anaesthesia induction and endoscopy.[31,32] **Predictors:** active escalation, current GI symptoms, higher dose. **Prevention:** follow anaesthesia-society guidance to hold the weekly dose around elective procedures, consider gastric ultrasound, prolong clear-fluid fasting, and treat as a full stomach if in doubt. Note the confounder: residual gastric contents are also more common in long-standing diabetic gastroparesis *without* any incretin agonist.

## 17.5 Loss of lean (muscle) mass

Because tirzepatide produces the largest weight loss, the absolute lean-mass component is also largest; the concern is functional sarcopenia in older patients. **Prevention** is behavioural: protein ≈ 1.2–1.6 g/kg/day and **progressive resistance training** throughout treatment (Chapter 18).[42]

## 17.6 Two feared risks that have been substantially de-risked

**Thyroid C-cell tumours.** The boxed contraindication derives from rodent C-cell hyperplasia; human population data have **not** shown an increase in medullary thyroid carcinoma. The contraindication remains precautionary in MTC/MEN-2 families.

**Suicidality / psychiatric.** Large analyses have **not** confirmed an increased risk of suicidal ideation; routine mood vigilance remains sensible but the population signal has been reassuring.[33]

## 17.7 Less common but important

Dehydration from GI losses can precipitate **acute kidney injury** (volume-mediated) — maintain hydration, especially during intercurrent illness. **Euglycaemic ketoacidosis** has been reported, particularly with concomitant SGLT2 inhibition or very low carbohydrate intake. Tirzepatide may modestly reduce the efficacy of **oral contraceptives** around initiation/escalation (delayed absorption) — counsel accordingly. For the eye, see Chapter 16.

## 17.8 A framework for weighing adverse effects

Because the adverse effects are numerous but mostly predictable, a framework serves better than a list. Three questions organise the chapter. First, *which of the two mechanisms does this effect arise from?* — the gastric-emptying/area-postrema pharmacology (nausea, vomiting, early satiety, aspiration risk) or the consequences of rapid weight loss (gallstones, part of the pancreatitis signal, lean-mass loss) — which tells you whether slow titration or gradual, supported weight loss is the relevant preventive lever. For tirzepatide, the pharmacological family carries extra weight because of its highest-in-class vomiting signal, and the loss-related family because it produces the largest weight loss, so both preventive levers matter more here than for a selective agonist. Second, *what is this patient's baseline risk of this event?* — a

18 PRACTICE

# Practical Management — Diet, Muscle, Plateau and Discontinuation

Tirzepatide is a tool that makes durable lifestyle change *possible*, not a substitute for it — and because it produces the *largest* weight loss of the licensed incretins (Chapter 12), the practical management of that loss matters more here than for any other agent. The quieting of hunger and food noise (Chapters 6 and 7) removes the biological headwind that defeats dietary effort, but the quality of the loss — above all the preservation of muscle — the handling of plateaus, and what happens on stopping all depend on how the treatment period is managed. This chapter translates the pharmacology into diet, training, monitoring and long-term strategy, with a heightened emphasis on lean-mass protection appropriate to the magnitude of loss the drug achieves.

## 18.1 Diet and protein

Because appetite is strongly suppressed, patients eat considerably less, making the *composition* of intake more important than the quantity. The first priority is **adequate protein**: during weight loss, sufficient dietary protein is the primary nutritional defence of lean mass, with a reasonable target of **1.2–1.6 g per kg per day** (using a sensible reference or goal weight in severe obesity), distributed across meals and favouring high-quality sources supplying the full amino-acid range. Because total intake is small and the weight loss is large, deliberately prioritising protein at each meal — and sometimes supplementing it — is often necessary to hit the target. The second priority is **nutrient density**: with fewer calories, the diet must be rich in vegetables, fruit, whole grains and micronutrient sources so that the reduced intake does not create nutritional gaps; adequate fibre eases constipation and hydration supports gastrointestinal comfort and renal protection during any vomiting or diarrhoea (Chapter 17). The measures that support nutrition also ease tolerability — smaller, lower-fat, unhurried meals — so good dietary practice and good tolerability reinforce one another, which matters especially given tirzepatide's higher gastrointestinal signal. Where intake is very restricted or prolonged, attention to specific micronutrients (including vitamin B12) is reasonable.

## 18.2 Muscle — resistance training is not optional

This is the most important practical message for tirzepatide specifically, because the chief downside of large weight loss scales with its magnitude, and tirzepatide produces the largest loss. Any substantial weight loss includes a component of **lean (muscle) mass**, and losing muscle carries real costs: reduced strength and mobility, a lower resting metabolic rate that predisposes to later regain, and — in older patients especially — the risk of clinically significant **sarcopenia** and frailty. With the ~20%+ losses tirzepatide can achieve, protecting muscle is not a refinement but a central goal, and the aim of good management is favourable *body recomposition* — losing fat while preserving or building muscle.

The intervention that preserves lean mass is **resistance (strength) training**, and for a drug of this potency it is genuinely not optional. A practical programme involves progressive

resistance exercise two to three times per week, working the major muscle groups through a full range of motion with a load and volume that increase gradually over time (progressive overload), begun early and continued throughout. It works synergistically with the protein intake of Section 18.1 — protein is the substrate, training the stimulus — so the two must be pursued together. Aerobic exercise has its own cardiovascular and metabolic benefits and should be encouraged, but it is *resistance* training that defends muscle and must not be displaced by cardio alone. For older patients the priority is greatest of all. Muscle-preserving pharmacological agents are in development (Chapter 19) and may eventually complement this, but for now protein plus resistance training is the established, effective strategy — and with tirzepatide it should be treated as an integral part of the prescription, not an afterthought.

## 18.3 The plateau

A plateau is an expected physiological equilibrium, not a failure, analysed in detail in Chapter 11. As weight falls, energy expenditure declines and the defended set-point partially reasserts itself, so loss ceases at a new equilibrium. The response depends on the dose: if below 15 mg and tolerating, escalation to the next step is the primary lever (having excluded a pseudo-plateau from incomplete titration or a dose not yet at steady state); if at 15 mg — the maximum, and the highest-efficacy licensed agent — the levers are behavioural and strategic: reinforce the diet and resistance-training programme, verify adherence, and reconsider whether the plateau is an acceptable maintenance weight, recognising that there is currently no higher-efficacy licensed agent to switch up to. The dosing lever and the lifestyle lever are used together.

## 18.4 Discontinuation and weight regain

The most important long-term counselling point, to be discussed *before* starting, is that discontinuation predictably leads to **weight regain**. **SURMOUNT-4** showed that stopping tirzepatide after successful loss led to regain of a large proportion of the lost weight, while continuation maintained it (Chapter 12). This is not a failure of will or drug but the direct consequence of the set-point physiology of Chapter 6: the drug overrides the defended set-point while present, and its removal lets the set-point reassert itself, restoring appetite and, with it, the weight, the ectopic fat and the metabolic risk. The implications: obesity should be framed and treated as a **chronic, relapsing disease** requiring long-term therapy, analogous to hypertension; this framing should shape the initial conversation and expectations so that continuation is planned rather than experienced as dependence; and where discontinuation is necessary or chosen, it should be accompanied by intensified lifestyle support and, ideally, a maintenance plan — with honesty that reliable "off-ramp" evidence (tapering to a lower dose or to lifestyle alone without regain) is currently limited (Chapter 11). Where the drug is continued, the maintenance goal is the lowest effective dose that holds the result with good tolerability, combined with the diet and training of this chapter.

## 18.5 Monitoring and the therapeutic relationship

Monitoring should track the whole integrated phenotype (Chapter 6), not weight alone: weight and, ideally, waist circumference and body composition to track fat versus lean change (particularly important given the large losses); HbA1c in diabetes; blood pressure and lipids;

19 LANDSCAPE

# The Incretin Landscape and Future Directions (2026)

Tirzepatide is the current efficacy leader among licensed incretins, and the proof-of-concept that *multi-receptor* agonism beats single-receptor agonism. This closing chapter places it in the landscape of 2026, surveys what is newer or coming, identifies the directions of travel and the open questions, and draws together the thesis that has run through the whole monograph.

## 19.1 Where tirzepatide sits

Tirzepatide produces greater weight loss than semaglutide head-to-head (SURMOUNT-5) and the greatest glycaemic effect of the class (SURPASS), and it holds a unique approval in obstructive sleep apnoea, with heart-failure (SUMMIT) and broader cardiovascular-outcome data maturing. Its position, in short, is that of the most *potent* licensed incretin. What it does not yet hold — a point made honestly in Chapters 13 and 14 — is the mature, placebo-controlled obesity cardiovascular- and renal-outcome evidence that semaglutide already has (SELECT, FLOW); those data, for tirzepatide, are awaited (SURMOUNT-MMO). Agent choice therefore balances tirzepatide's magnitude and its unique OSA/HFpEF indications against tolerability (its highest-in-class gastrointestinal signal), cost, the injectable-only route, and the specific outcome evidence a given patient needs.

## 19.2 Multi-receptor agonism is the future

Tirzepatide proved the central principle now driving the field: that recruiting additional hormonal pathways beyond GLP-1 yields greater efficacy. The extension of that principle is already in late development. The triple GIP/GLP-1/glucagon agonist **retatrutide** produced weight loss of around 24% in phase 2 and is in phase 3 (the TRIUMPH programme), with projected availability toward 2028. Its innovation is the addition of a *glucagon-receptor* arm on top of the GIP and GLP-1 arms: although glucagon raises hepatic glucose output (undesirable in diabetes), it also increases energy expenditure and promotes hepatic fat oxidation (desirable for weight and fatty liver), and the triple agonist exploits those benefits while relying on its potent GLP-1 and GIP activity to keep glycaemia controlled (Chapter 4). If retatrutide fulfils its early promise, it will exceed tirzepatide as tirzepatide exceeded the selective agonists — extending medical therapy into the weight-loss range previously reached only by surgery. The dual-agonist principle tirzepatide established is, in this sense, a waypoint on a trajectory toward ever-greater efficacy through complementary multi-receptor pharmacology.

## 19.3 Oral and small-molecule agents

The other defining trend is the move away from injection. An *oral tirzepatide* is in development but not yet approved, so the dual agonist remains injectable for now (Chapter 8). Meanwhile the first oral *non-peptide, small-molecule* GLP-1 receptor agonist, **orforglipron**, reached the market in 2026 for weight management. Because it is a small molecule rather than a peptide, orforglipron is far easier and cheaper to manufacture at scale and needs

neither an absorption-enhancer nor a cold chain — it is a conventional pill. Its weight-loss magnitude (around 12%) is a *single-incretin* effect, below tirzepatide's, but its scalability and convenience could greatly widen access. The clear longer-term trajectory is toward agents that combine the needle-free convenience of a small-molecule oral with the higher efficacy of multi-receptor agonism — the best of both directions.

## 19.4 Other mechanisms — amylin and muscle preservation

Two further directions matter, one of them especially relevant to tirzepatide's profile. The first is the **amylin** axis: amylin analogues (such as cagrilintide, including in combination with semaglutide as "CagriSema") act through a distinct satiety pathway and, in combination with incretin agonism, produce high degrees of weight loss — another hormonal lever being added to the field. The second, and particularly pertinent given tirzepatide's class-leading weight loss, is the development of **muscle-preserving agents** (for example, therapies targeting the myostatin/activin pathway) intended to preserve or build lean mass during weight loss. Because the concern about lean-mass loss scales with the magnitude of weight loss (Chapter 18), such agents could be especially valuable in combination with the most potent incretins like tirzepatide, improving the *quality* of the loss — more fat, less muscle — rather than only its magnitude. If they prove safe and effective, they could complement the protein-and-resistance-training strategy that currently addresses lean-mass loss behaviourally.

## 19.5 Access, indications, and open questions

As the tirzepatide shortage has resolved and the FDA moves to close bulk compounding (Chapter 11), the cheap-compounded era is narrowing, shifting prescribing back toward labelled product and making *cost and access*, rather than pharmacology, the dominant near-term constraint on who benefits from therapies this effective. Indications continue to expand — the OSA approval opened the conceptual door to treating the mechanical diseases of adiposity pharmacologically, and heart-failure and other indications are maturing. Important questions remain open: the durability problem and whether a reliable "off-ramp" from continuous therapy is possible (Chapters 12 and 18); how best to preserve muscle during the largest weight losses; the maturing cardiovascular and renal outcome evidence (SURMOUNT-MMO and a dedicated renal trial); the role of these agents in children and adolescents; long-term safety over decades of use; and the profound equity and cost issues that determine access. The resolution of these questions will shape the field as much as the new molecules.

## 19.6 Closing synthesis

It is worth ending where the monograph began. Tirzepatide is the first clinical vindication of an idea that had long been doubted — that GIP, dismissed as a therapeutic dead end because it failed in isolation in the diabetic beta cell, could become a powerful partner when combined with GLP-1 in a single, carefully tuned molecule. From that insight, and the imbalanced, biased dual-receptor pharmacology that realises it, flows everything this monograph has developed: the convergence of two incretin signals on glucose-dependent insulin secretion; the net glucagon suppression that restrains the liver; the distinctive GIP adipose arm that improves insulin sensitivity beyond what weight loss alone predicts; the central appetite and reward effects, with the GIP contribution to tolerability; and the class-leading efficacy that has