A CLINICAL MONOGRAPH · CONTINUING MEDICAL EDUCATION



Semaglutide

The Long-Acting GLP-1 Receptor Agonist — Mechanism of Action, Whole-Body Effects, Patient Selection, Dosing and Safety

A complete mechanism-to-bedside reference

How semaglutide works at the receptor, the islet, the liver, muscle, the brain's appetite and reward circuits and the vasculature — and what that mechanism delivers in glycaemia, weight, cardiovascular, renal, hepatic, ophthalmic and addiction outcomes — followed by patient selection, full dosing (subcutaneous, oral, 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 Incretin Physiology to Semaglutide

2 Molecular Structure and Oral Delivery

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, Ectopic Fat 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 — Subcutaneous and Oral

9 Patient Selection — For Whom, and For Whom Not

10 Dosing I — Subcutaneous and Oral Titration

11 Dosing II — Microdosing, Macrodosing and the Current Landscape

12 Glycaemic and Weight Efficacy — STEP, SUSTAIN and PIONEER

13 Cardiovascular Outcomes and Their Mechanism — SELECT and SUSTAIN-6

14 Renal Outcomes and Their Mechanism — FLOW

15 Liver, Heart Failure and Sleep Apnoea — ESSENCE, STEP-HFpEF and Beyond

16 The Eye — Diabetic Retinopathy and the NAION Signal

17 Adverse Effects — Mechanisms, Predictors, Prevention and Confounders

18 Practical Management — Diet, Muscle, Plateau and Discontinuation

19 The GLP-1 Landscape and Future Directions (2026)

R References

How to read this monograph

This edition is built to make the reader fluent in semaglutide from molecule to bedside. Chapters 3–7 develop the pharmacology in depth — receptor signalling, suppression of hepatic glucose production, the determinants of insulin sensitivity, central appetite control, 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** each benefit occurs. 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 Incretin Physiology to Semaglutide

Semaglutide is best understood not as a novel chemical curiosity but as the most refined expression of a physiological idea more than a century old: that the gut, on sensing nutrients, sends hormonal signals to the pancreas that prepare the body to handle the incoming glucose load. Almost every clinically relevant property of the drug — its glucose-dependent action, its once-weekly dosing interval, its gastrointestinal side-effect profile, its appetite and reward effects, even the strict rules around its oral formulation — follows directly from deliberate molecular engineering designed to overcome the limitations of the native hormone it imitates. To prescribe semaglutide intelligently, therefore, one must first understand the system it exploits. This opening chapter develops that system in depth: the incretin effect and its discovery, the two incretin hormones and how they are made and released, the full repertoire of GLP-1's actions, the way the system fails in type 2 diabetes, the pharmacological obstacle that kept native GLP-1 from ever becoming a drug, and the engineering arc that culminated in a once-weekly and ultimately oral agonist. The remaining chapters build on this foundation, but the logic of the entire monograph is set here.

1.1 The incretin effect — a century-old observation

The conceptual seed was planted in the early twentieth century. As early as 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** — a contraction of "INtestinal seCRETion of INsulin" — for a putative gut substance that lowered blood glucose by promoting internal (endocrine) pancreatic secretion, distinguishing it from the digestive "excretin" fraction. For decades the idea remained unproven, because the tools to measure circulating insulin did not yet exist.

That changed with the radioimmunoassay for insulin developed by Berson and Yalow in the late 1950s, which for the first time allowed insulin to be quantified in plasma. 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 a substantially greater insulin response 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 not a minor adjustment but 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; the larger the glucose load, the greater the proportional contribution. In other words, for most of a normal day's meals, the majority of the insulin response is being driven not by glucose acting directly on the beta cell but by hormones the gut releases in anticipation of, and during, nutrient absorption. This single fact reframes the beta cell as the downstream

executor of a gut-initiated programme, and it explains why a drug that mimics those gut signals can have such a large and physiological effect on glucose handling.

1.2 Two incretin hormones — GIP and GLP-1

Once the incretin effect was firmly established, the search turned to the molecules responsible. Two were ultimately identified, and together they account for essentially the entire incretin effect in humans.

The first to be characterised was **glucose-dependent insulinotropic polypeptide** (GIP), originally named gastric inhibitory polypeptide for its early-described ability to inhibit gastric acid secretion at supraphysiological doses. GIP is a 42-amino-acid peptide secreted by the enteroendocrine **K cells**, which are concentrated in the proximal small intestine — the duodenum and proximal jejunum. As nutrients (particularly fat and glucose) reach this region, K cells release GIP into the circulation, where it acts on its receptor on the pancreatic beta cell to augment glucose-stimulated insulin secretion. GIP is the quantitatively larger contributor to the incretin effect in health, but, as discussed below, its insulinotropic action is markedly blunted in type 2 diabetes, which limited its appeal as a therapeutic target for many years (an assessment that the success of the dual agonist tirzepatide has since complicated).

The second incretin, and the one this monograph is principally concerned with, is **glucagon-like peptide-1** (GLP-1). GLP-1 is produced by the enteroendocrine **L cells**, which, in contrast to K cells, are distributed predominantly in the distal small intestine and colon, although they are present throughout the gut. GLP-1 derives from the **proglucagon** gene, a single gene whose primary transcript is processed in a tissue-specific manner by different prohormone convertase enzymes. In the pancreatic alpha cell, proglucagon is cleaved by prohormone convertase 2 to yield glucagon itself. In the intestinal L cell and in certain neurons of the brainstem, the same proglucagon precursor is processed instead by prohormone convertase 1/3 to yield GLP-1, GLP-2, oxyntomodulin and glicentin. This is an elegant example of post-translational versatility: one gene, several biologically distinct peptides, determined entirely by which convertase is expressed in a given cell. It also explains an important pharmacological point — that GLP-1 and glucagon share a common ancestor and considerable structural homology, which is why some engineered agonists (the glucagon/GLP-1 co-agonists and triple agonists discussed in Chapter 19) can engage both receptor families.

Native GLP-1 circulates in two principal bioactive forms, GLP-1(7-37) and the more abundant amidated GLP-1(7-36)amide, the numbering reflecting cleavage of the proglucagon-derived sequence. Both bind the GLP-1 receptor with high affinity and are, for practical purposes, equivalent in their actions. The therapeutic analogues, including semaglutide, are built upon this GLP-1(7-37) backbone.

1.3 How GLP-1 is secreted

Understanding the triggers for L-cell secretion is more than academic, because it clarifies why incretin-based therapy is fundamentally a meal-related, nutrient-sensing system. L cells are open-type enteroendocrine cells: they possess apical microvilli that face the gut lumen and can directly sense its contents, while their basolateral surface abuts the capillary network into which they secrete. They respond to all three macronutrient classes. Carbohydrate, particularly glucose, is sensed in part through the sodium-glucose cotransporter SGLT1,

02 CHEMISTRY

Molecular Structure and Oral Delivery

If Chapter 1 explained *why* a long-acting GLP-1 agonist was needed, this chapter explains *how* the molecule was built to be one. Semaglutide is not a found compound but a deliberately engineered one, and every atom that distinguishes it from native GLP-1 was placed to solve a specific problem: resistance to enzymatic degradation, prolongation of the dosing interval to once weekly, and — most remarkably — the achievement of meaningful absorption when swallowed, a feat that had eluded peptide pharmacology for a century. A clinician who understands these three structural solutions can predict almost everything about how the drug behaves: why it lasts a week, why a missed dose can be taken days late, why the oral tablet must be taken on an empty stomach with almost no water, why effects persist for weeks after stopping, and why the compounded "grey-market" versions that proliferated during the shortage years carried real risk. This chapter develops the chemistry in the depth those clinical inferences require.

2.1 The design brief and the GLP-1(7-37) scaffold

Semaglutide is a **GLP-1 receptor agonist** built on the backbone of human GLP-1(7-37), the 31-amino-acid bioactive form of the native incretin. Retaining the native sequence as the scaffold was a deliberate choice: it preserves full agonist activity at the human GLP-1 receptor and keeps the molecule maximally "self," minimising the immunogenicity that plagued earlier, more exotic peptide drugs. The design brief was then to make three, and only three, targeted changes to this scaffold — changes calculated to confer the desired pharmacokinetic properties while leaving the receptor-binding pharmacophore intact.

The discipline of changing as little as possible is itself an important principle of peptide drug design. Each substitution risks altering receptor affinity, signalling bias, or immunogenicity; the art lies in achieving the pharmacokinetic goal with the fewest and most conservative modifications. Semaglutide's developers achieved a roughly one-week half-life and full agonism with just three modifications to the native sequence, a parsimony that is part of why the molecule is so well tolerated immunologically and so reliable pharmacologically. The three changes are summarised in Table 2.1 and then dissected individually.

Table 2.1. The three structural modifications of semaglutide and their purpose.

Modification	Location	Purpose
Aib (α-aminoisobutyric acid) substitution	Position 8 (replacing Ala8)	Blocks DPP-4 cleavage at the N-terminus
C18 fatty diacid + γGlu/2×OEG linker	Acylated to Lys at position 26	Reversible albumin binding → ~1-week half-life
Arg substitution	Position 34 (replacing Lys34)	Directs acylation to the single correct lysine

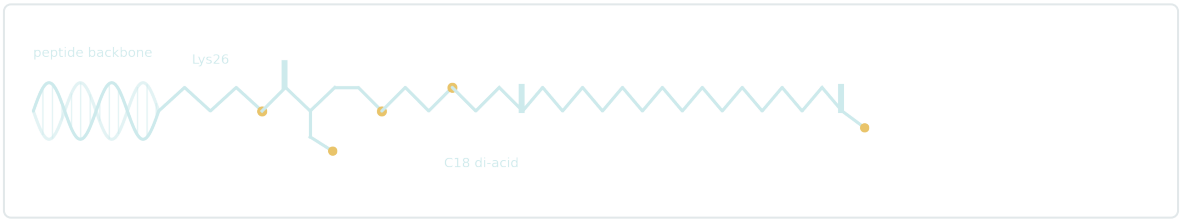


Figure 2.1. Schematic of semaglutide: the GLP-1 peptide backbone (helical ribbon) bearing, at Lys26, the acyl side chain — a γ-glutamate spacer and two OEG units (oxygen atoms marked) linking the peptide to a C18 fatty di-acid that binds albumin.

2.2 The Aib8 substitution — defeating DPP-4

The first and most fundamental modification addresses the enzyme that destroys native GLP-1. As established in Chapter 1, DPP-4 cleaves the N-terminal His-Ala dipeptide from GLP-1(7-37) because position 8 (the second residue of the bioactive peptide) is an alanine, exactly the residue DPP-4 recognises. The structural solution is to replace this alanine with **α-aminoisobutyric acid** (Aib), a non-proteinogenic amino acid.

Why Aib in particular? Aib is a doubly methylated, sterically hindered amino acid — essentially alanine with an extra methyl group on the α-carbon, making that carbon quaternary. This bulk does two things. First, it sterically obstructs the active site of DPP-4, preventing the enzyme from engaging and cleaving the peptide bond; the N-terminus is thereby protected and the peptide retains its intact, active form. Second, Aib has a strong conformational influence: its restricted geometry stabilises the helical structure of the N-terminal region, which happens to be important for high-affinity binding and activation of the GLP-1 receptor. So the same substitution that confers protease resistance also subtly favours the bioactive conformation — a fortunate convergence. The use of Aib at this position is now a hallmark of incretin drug design and recurs in the dual agonist tirzepatide.

The clinical consequence of the Aib8 substitution is profound: it converts a peptide with a 1–2 minute half-life (from DPP-4 cleavage) into one whose degradation by that route is essentially eliminated. DPP-4 resistance alone, however, is not sufficient for once-weekly dosing, because the small, intact peptide would still be cleared rapidly by the kidney. That second problem required the acylation strategy.

2.3 The acylation — a C18 fatty diacid and its linker

The second and signature modification is the attachment of a lipid chain that confers reversible albumin binding, the protraction strategy introduced in Chapter 1. In semaglutide this is realised with precision. A **C18 fatty diacid** (an eighteen-carbon chain bearing a carboxylic acid at each end — a stearic diacid) is attached, not directly, but through a carefully engineered **linker** to the side-chain amino group of a lysine at position 26 of the peptide.

The linker is itself a deliberate piece of chemistry, composed of a γ-glutamic acid (γGlu) spacer and two units of a short polyethylene-glycol-like moiety known as **OEG** (8-amino-3,6-dioxaoctanoic acid). Reading from the peptide outward, the architecture is: Lys26 side chain → γGlu → OEG → OEG → C18 diacid. Each component serves a purpose. The γGlu and OEG spacers provide length and flexibility, holding the fatty acid far enough from the peptide that it can insert freely into albumin's binding pocket without the peptide's own bulk getting in the way, and improving the solubility and the kinetics of binding. The terminal carboxylic acid of the C18 diacid is important too: the dicarboxylic (rather than monocarboxylic) design

03 MECHANISM I

Receptor Signalling and Glucose-Dependent Insulin Secretion

This is the first of the five mechanism chapters, and in some ways the most important, because it establishes the property that defines the safety of the entire drug class: that semaglutide enhances insulin secretion *only when glucose is elevated*. A clinician who understands this single fact at a mechanistic level can explain why the drug lowers glucose powerfully yet almost never drives it below normal on its own; why hypoglycaemia, when it occurs, is the fault of a co-prescribed sulfonylurea or insulin rather than of semaglutide; why the drug works best in patients who retain some beta-cell function; and why the glucose-lowering effect is durable rather than self-exhausting. To reach that understanding, this chapter first reconstructs how the healthy beta cell senses glucose and releases insulin, then shows precisely where and how GLP-1-receptor signalling intervenes — and, crucially, why that intervention is gated by glucose.

3.1 The beta cell as a glucose sensor

Insulin secretion from the pancreatic beta cell is, at its core, a mechanism for translating the ambient glucose concentration into a proportional amount of secreted insulin. The elegance of the system lies in how directly the cell couples glucose metabolism to electrical activity and then to secretion. This is conventionally called the **triggering pathway**, and it proceeds in an ordered sequence.

Glucose enters the beta cell through facilitative glucose transporters down its concentration gradient, so that intracellular glucose tracks the blood level. Once inside, glucose is phosphorylated by the enzyme **glucokinase**, and this step is the true "glucose sensor" of the cell. Glucokinase is distinctive among hexokinases: it has a relatively high Michaelis constant, meaning its activity rises steeply across exactly the physiological range of blood glucose, and it is not inhibited by its product. Because the rate of glucose phosphorylation — and therefore the rate of everything downstream — is set by glucokinase in proportion to the glucose concentration, glucokinase determines the threshold and the sensitivity of insulin secretion. It is, in effect, the cell's glucose meter.

Phosphorylated glucose then enters glycolysis and, ultimately, mitochondrial oxidative metabolism, generating ATP. The consequence is a rise in the intracellular **ATP-to-ADP ratio** in proportion to the glucose load. This rising ATP/ADP ratio is the metabolic signal that the cell converts into an electrical one, through the next, pivotal component.

That component is the **ATP-sensitive potassium channel** (K_{ATP} channel), an octameric channel built from four pore-forming Kir6.2 subunits and four regulatory SUR1 (sulfonylurea receptor) subunits. At low glucose, when ATP/ADP is low, these channels are open, allowing potassium to flow out of the cell and holding the membrane at a negative (hyperpolarised, resting) potential. As glucose rises and ATP/ADP increases, ATP binds the Kir6.2 subunits and *closes* the K_{ATP} channels. The outward potassium leak ceases, and the membrane **depolarises**. This is the master switch of the system: the K_{ATP} channel converts the metabolic glucose signal into a change in membrane voltage. It is also, not coincidentally, the target of

sulfonylureas, which close the channel directly through the SUR1 subunit regardless of glucose — a fact central to understanding hypoglycaemia risk later in this chapter.

Depolarisation opens voltage-gated calcium channels in the plasma membrane, and calcium floods into the cell down its steep electrochemical gradient. The resulting rise in intracellular calcium is the immediate trigger for the final step: the calcium-dependent **exocytosis** of insulin-containing secretory granules, which fuse with the plasma membrane and release insulin into the circulation. The chain is therefore: glucose → glucokinase → ATP/ADP rise → K_{ATP} closure → depolarisation → calcium influx → insulin exocytosis. Each step is proportional to glucose, so the amount of insulin secreted tracks the glucose concentration — exactly the behaviour a glucose-regulated hormone requires.

It is worth noting that insulin secretion is biphasic. A brief, sharp **first phase** reflects the release of a small pool of granules already docked and primed at the membrane, discharged within minutes of a glucose rise; a more sustained **second phase** follows as further granules are mobilised, primed and released. The loss of first-phase insulin secretion is one of the earliest detectable beta-cell abnormalities in type 2 diabetes, and, as will be seen, GLP-1 signalling acts on precisely the machinery that governs both phases.

3.2 From the GLP-1 receptor to cyclic AMP

Against this background, the action of GLP-1 becomes intelligible. The GLP-1 receptor is a **class B G-protein-coupled receptor** expressed on the beta-cell surface, coupled to the stimulatory G-protein **G_s**. When semaglutide binds and activates the receptor (engaging it through the two-domain mechanism described in Chapter 2), G_s in turn activates the membrane enzyme **adenylyl cyclase**. Adenylyl cyclase synthesises the second messenger **cyclic AMP (cAMP)** from ATP, and intracellular cAMP rises. cAMP is the central node of GLP-1 action in the beta cell: essentially all of the acute insulinotropic effects of GLP-1 flow from this rise in cAMP. The receptor itself does nothing to glucose; it raises cAMP, and cAMP then modulates the glucose-sensing and secretory machinery described above.

3.3 The two arms of cAMP action — PKA and Epac2

Cyclic AMP exerts its effects through two principal downstream effectors, and appreciating both is necessary to understand the full reach of GLP-1 signalling. The first is **protein kinase A** (PKA), the classical cAMP-dependent kinase. When cAMP binds PKA's regulatory subunits, the catalytic subunits are freed and phosphorylate a wide range of target proteins — ion channels, components of the exocytotic machinery, and transcription factors. The second effector is **Epac2** (exchange protein directly activated by cAMP 2, also called cAMP-GEFII), a cAMP-binding protein that acts independently of PKA, principally on the secretory machinery and intracellular calcium handling. The existence of two parallel effectors — one a kinase acting broadly, one a guanine-nucleotide exchange factor acting more locally on exocytosis and calcium — is why GLP-1's amplifying effect touches several distinct steps of insulin secretion rather than a single one. PKA tends to dominate the longer-range and transcriptional effects; Epac2 contributes strongly to the rapid potentiation of granule exocytosis. Together they constitute the molecular machinery through which a rise in cAMP is translated into more insulin.

04 MECHANISM II

Glucagon Suppression and the Control of Gluconeogenesis

Chapter 3 explained how semaglutide 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 — and it 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 not in too little insulin after meals but in too much glucose pouring out of the liver around the clock. To understand how semaglutide corrects this, one must understand the biochemistry of how the liver makes glucose, how the hormone glucagon switches that machinery on, how insulin switches it off, and how the drug shifts the balance between them. This is among the most mechanistically detailed chapters in the monograph, because the request that prompted it — to explain "how gluconeogenesis is happening and how it is being stopped by increasing insulin" — is answered precisely here, at the level of the enzymes and transcription factors involved.

4.1 Hepatic glucose production and fasting hyperglycaemia

The liver is the body's glucose factory and reservoir. Between meals and overnight, when no glucose is being absorbed from the gut, it is the liver that maintains blood glucose by releasing glucose into the circulation — a process essential for supplying the brain, which depends on a continuous glucose supply. The liver does this by two routes. The first is **glycogenolysis**, the breakdown of stored glycogen, which dominates in the early hours of fasting but is limited by the finite size of the glycogen store. The second, which sustains glucose production through more prolonged fasting and which is the central concern of this chapter, is **gluconeogenesis** — the synthesis of new glucose molecules from non-carbohydrate precursors.

In health, hepatic glucose production is exquisitely regulated: it is switched on during fasting and switched off after meals, when insulin rises and signals abundance. In type 2 diabetes this regulation breaks down. Hepatic glucose production becomes excessive and is no longer appropriately suppressed in the fed state, so the liver continues to pour 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 — the elevated morning glucose that patients and clinicians see, and a major contributor to overall glycaemic burden. Any therapy that can rein in hepatic glucose production therefore addresses a core mechanism of the disease, and this is exactly what the glucagon-suppressing action of semaglutide accomplishes.

4.2 The biochemistry of gluconeogenesis

To appreciate how the process is controlled, it is necessary to know its outline. Gluconeogenesis builds a six-carbon glucose molecule from smaller precursors, principally **lactate** (from anaerobic glycolysis in muscle and red cells, returned to the liver via the Cori cycle), **glycerol** (released from adipose triglyceride during lipolysis), and **glucogenic amino**

acids, especially alanine (delivered from muscle protein breakdown via the glucose–alanine cycle). The pathway is, in broad terms, glycolysis run in reverse, but three steps of glycolysis are thermodynamically irreversible and must be bypassed by distinct gluconeogenic enzymes. These bypass enzymes are the regulated control points of the whole pathway, and it is precisely these that glucagon induces and insulin represses, so they are worth naming.

The first bypass converts pyruvate to phosphoenolpyruvate in two steps: **pyruvate carboxylase** (in the mitochondrion) produces oxaloacetate, and **phosphoenolpyruvate carboxykinase** (**PEPCK**, gene *PCK1*) then produces phosphoenolpyruvate. PEPCK is a key rate-controlling, hormonally regulated enzyme. The second bypass is catalysed by **fructose-1,6-bisphosphatase** (FBPase-1), which removes a phosphate from fructose-1,6-bisphosphate; this enzyme is the reciprocal counterpart of the glycolytic enzyme phosphofructokinase-1 (PFK-1), and the two together form a critical regulatory switch (discussed below). The third bypass is catalysed by **glucose-6-phosphatase** (**G6Pase**, gene *G6PC*), which removes the final phosphate to yield free glucose that can leave the hepatocyte. PEPCK and G6Pase are the two enzymes whose names recur throughout this chapter, because their transcription is the principal long-term lever by which glucagon and insulin control gluconeogenic capacity.

A crucial regulatory node sits at the fructose-1,6-bisphosphatase / phosphofructokinase step, governed by a signalling metabolite called **fructose-2,6-bisphosphate**. Fructose-2,6-bisphosphate is a potent *activator* of glycolysis (it stimulates PFK-1) and a potent *inhibitor* of gluconeogenesis (it inhibits FBPase-1). Its level is set by a single bifunctional enzyme (PFK-2/FBPase-2) whose activity is controlled by phosphorylation. When fructose-2,6-bisphosphate is high, the liver burns glucose (glycolysis); when it is low, the liver makes glucose (gluconeogenesis). This single metabolite is therefore a master switch determining the direction of hepatic carbohydrate flux, and, as the next section shows, glucagon lowers it while insulin raises it — which is how the two hormones flip the liver between glucose production and glucose consumption.

4.3 How glucagon drives gluconeogenesis

Glucagon, secreted by the pancreatic alpha cell during fasting and falling glucose, is the hormone that commands the liver to produce glucose. It binds the hepatocyte **glucagon receptor**, a G_s-coupled class B GPCR (a relative of the GLP-1 receptor), raising cAMP and activating **protein kinase A** (PKA). PKA then drives hepatic glucose output through three coordinated actions, summarised in Table 4.1 and elaborated here.

First, PKA stimulates **glycogenolysis**: through a phosphorylation cascade (PKA activates phosphorylase kinase, which activates glycogen phosphorylase) it accelerates the breakdown of glycogen to glucose, and simultaneously inhibits glycogen synthase, halting glycogen storage. This provides the rapid, early-fasting glucose release. Second, and central to this chapter, PKA **induces the gluconeogenic genes**: it phosphorylates and activates the transcription factor CREB, which, together with the transcriptional coactivator PGC-1α and the transcription factor FOXO1, drives increased transcription of **PEPCK** and **glucose-6-phosphatase**. This raises the liver's gluconeogenic capacity — more enzyme, more glucose synthesis. Third, PKA **redirects metabolic flux toward glucose production**: by phosphorylating the bifunctional PFK-2/FBPase-2 enzyme, it lowers **fructose-2,6-bisphosphate**, which relieves the inhibition of FBPase-1 and removes the activation of PFK-1

05 MECHANISM III

Insulin Resistance, Ectopic Fat 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* — the failure of muscle, liver and fat to respond properly to insulin. This chapter explains what insulin resistance is at the molecular level, why ectopic fat is its dominant cause, how chronic inflammation contributes, and — most importantly for the prescriber — how the weight loss that semaglutide produces reverses it. This is where the drug's appetite and weight effects (Chapters 6 and 7) connect to its metabolic benefit: semaglutide improves insulin sensitivity substantially, and it does so principally by removing the fat that caused the resistance in the first place. Understanding this chapter explains why the magnitude of metabolic benefit tracks the magnitude of weight loss, why the drug improves not only glucose but blood pressure, lipids, fatty liver and inflammation together, and why the benefits form a self-reinforcing cycle.

5.1 What insulin does at the cell — the signalling cascade

To understand resistance, one must first understand normal insulin action. Insulin binds the insulin receptor, a receptor tyrosine kinase, on the surface of its target cells. Binding activates the receptor's intrinsic kinase, which phosphorylates the receptor itself and a family of docking proteins called **insulin receptor substrates** (IRS-1 and IRS-2) on specific *tyrosine* residues. Tyrosine-phosphorylated IRS proteins then recruit and activate the enzyme **PI3-kinase (PI3K)**, which generates the membrane lipid PIP_3, which in turn activates the central node of insulin's metabolic actions, the kinase **Akt** (also called protein kinase B).

Activated Akt then executes insulin's metabolic programme in a tissue-specific way. In **muscle and adipose tissue**, Akt drives the translocation of the glucose transporter **GLUT4** from intracellular storage vesicles to the cell surface, allowing glucose to enter the cell — this is how insulin clears glucose from the blood into muscle and fat, and skeletal muscle is quantitatively the largest site of insulin-stimulated glucose disposal. In the **liver**, as developed in Chapter 4, Akt phosphorylates and excludes FOXO1 from the nucleus, switching off gluconeogenesis, and inhibits GSK-3 to promote glycogen synthesis. Across tissues, Akt also suppresses lipolysis in adipose tissue and promotes protein and lipid synthesis. The IRS-1 → PI3K → Akt → GLUT4 axis in muscle, and the parallel hepatic axis, are the pathways whose failure constitutes insulin resistance.

5.2 What insulin resistance is, molecularly

Insulin resistance is, at the molecular level, a defect in this signalling cascade — the cells receive the insulin signal but transmit it poorly, so that a normal or even elevated insulin concentration produces a subnormal biological response. The body compensates by secreting more insulin (hyperinsulinaemia), which masks the defect for a time, until the beta cell can no longer keep pace and frank hyperglycaemia emerges.

The crucial molecular lesion lies at the level of IRS-1. In the healthy state, insulin drives *tyrosine* phosphorylation of IRS-1, which propagates the signal to PI3K. In the insulin-resistant state, IRS-1 instead carries excessive **serine phosphorylation** at inhibitory sites. Serine phosphorylation of IRS-1 impairs its ability to be tyrosine-phosphorylated by the insulin receptor and to engage PI3K; it effectively *uncouples* the receptor from its downstream signalling. With IRS-1 uncoupled, PI3K and Akt are inadequately activated, GLUT4 translocation fails, and glucose uptake into muscle falls — the hallmark of peripheral insulin resistance. The central question of insulin resistance therefore becomes: what causes this excessive, inhibitory serine phosphorylation of IRS-1? The answer, in most cases, is ectopic fat.

5.3 The dominant driver — ectopic fat and lipotoxicity

The single most important and best-characterised cause of insulin resistance is the accumulation of fat in tissues where fat does not belong — **ectopic fat** — and the toxic effects it exerts there, termed **lipotoxicity**. The logic runs as follows. Adipose tissue is the body's designated fat store, evolved to sequester surplus energy safely as triglyceride. When energy intake chronically exceeds expenditure, however, the storage capacity of (particularly subcutaneous) adipose tissue can be exceeded — the **adipose expandability** or "overflow" hypothesis — and lipid begins to deposit in tissues not designed to store it: skeletal muscle, liver and the pancreas. It is this ectopic lipid, far more than total body fat per se, that drives insulin resistance, which is why the *distribution* of fat (visceral and ectopic versus subcutaneous) matters more than weight alone.

The mechanistic chain by which ectopic lipid causes insulin resistance is now well characterised, and it is worth following in detail because it is the precise reason weight loss works. Excess intracellular lipid raises the level of **diacylglycerol (DAG)**, a lipid intermediate. DAG is a physiological activator of a family of enzymes called the **novel protein kinase C** isoforms — specifically PKCθ in skeletal muscle and PKCε in the liver. Activated by the accumulating DAG, these kinases phosphorylate IRS-1 (and the insulin receptor) on the inhibitory *serine* residues described above. The consequence is exactly the lesion of Section 5.2: serine-phosphorylated IRS-1 is uncoupled from the insulin receptor, PI3K/Akt signalling fails, and — in muscle — GLUT4-mediated glucose uptake falls, while — in liver — the failure to exclude FOXO1 leaves gluconeogenesis unrestrained (linking directly to Chapter 4). A parallel class of lipid intermediates, the **ceramides**, contributes through related mechanisms that also impair Akt activation. The essential point is that ectopic fat is not metabolically 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 — which is what makes weight loss mechanistically curative rather than merely palliative.

06 MECHANISM IV

Appetite, Gastric Emptying and the Integrated Phenotype

Weight loss — the engine of much of semaglutide's metabolic benefit (Chapter 5) — is driven mainly by central and gastrointestinal actions: a reduction in hunger, an earlier sense of fullness, a quieting of preoccupation with food, and a slowing of gastric emptying. This chapter develops these appetite mechanisms in depth, distinguishing the homeostatic control of hunger (the subject here) from the hedonic, reward-based control of craving (Chapter 7), and then steps back to show how the four mechanisms of Chapters 3–6 converge on a single, coherent cardiometabolic phenotype. Understanding this chapter explains why patients eat less without consciously dieting, why the gastrointestinal side effects and the satiety effect share an origin, and why a drug acting on appetite simultaneously improves glucose, blood pressure, lipids, hepatic fat and inflammation.

6.1 Two systems of appetite control

Food intake is governed by two intertwined systems. The **homeostatic** system regulates energy balance over time, matching intake to the body's energy needs and defending the body-weight set-point through the gut and adipose hormones described in Chapter 1; it is the system that generates and relieves hunger. The **hedonic** (reward) system, by contrast, governs the desire for the pleasure and reinforcement that palatable food provides, independent of energy need, and is the seat of craving; it is developed separately in Chapter 7. GLP-1 acts on both, but this chapter concerns the homeostatic arm — how semaglutide reduces hunger and increases satiety — while keeping in mind that the two systems are anatomically and functionally interwoven, so the distinction is one of emphasis rather than a clean separation.

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

The central hub of homeostatic appetite control is the **arcuate nucleus** of the hypothalamus, which contains two opposing populations of first-order neurons. One population expresses **POMC** (pro-opiomelanocortin) and CART; these are *anorexigenic* — when activated, they suppress appetite and promote satiety, largely by releasing α-melanocyte-stimulating hormone, which activates downstream melanocortin (MC4) receptors that reduce food intake and increase energy expenditure. The opposing population expresses **AgRP** (agouti-related peptide) and **NPY** (neuropeptide Y); these are *orexigenic* — when activated, they powerfully stimulate hunger and food-seeking, both by releasing NPY and by antagonising the melanocortin system. The balance of activity between these two populations sets the level of hunger at any moment, and they integrate hormonal signals of energy status (leptin, insulin, ghrelin, and the incretins) to do so.

GLP-1-receptor signalling shifts this balance toward satiety. Activation of GLP-1 receptors in and around the arcuate nucleus *stimulates the anorexigenic POMC/CART neurons* and *inhibits the orexigenic AgRP/NPY neurons*. The net effect is a reduction in hunger and an earlier and

stronger sense of fullness for a given amount of food. This is the homeostatic basis of reduced energy intake on semaglutide: not willpower, but a pharmacological resetting of the hunger–satiety balance toward satiety. Because this acts on the same melanocortin circuitry that defends the body-weight set-point (Chapter 1), it helps explain how the drug can lower the defended set-point rather than merely imposing a transient restriction that the homeostatic system fights to overturn.

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

The hypothalamus is not the only site of appetite control engaged by semaglutide. The **hindbrain** — particularly the nucleus of the solitary tract (NTS) and the adjacent **area postrema** — is a second major centre, integrating signals about satiety, nausea and visceral state. These structures are critical to understanding how a large, peripherally administered, albumin-bound peptide exerts central effects despite the blood–brain barrier. The area postrema is a **circumventricular organ**: one of a small number of brain regions where the blood–brain barrier is incomplete, allowing circulating peptides direct access to neural tissue. In addition, peripheral GLP-1 signals reach the NTS through **vagal afferent** nerves carrying information from the gut. Through these two routes — direct access at the area postrema and indirect signalling via the vagus — semaglutide engages the hindbrain appetite circuitry that then communicates with the hypothalamus and higher centres.

A point of real clinical importance is that the area postrema is also the brain's principal **emetic sensor** — the chemoreceptor trigger zone that detects circulating substances and can initiate nausea and vomiting. The very accessibility that allows semaglutide to engage hindbrain satiety circuitry is therefore also what underlies the nausea that is the drug's commonest adverse effect (Chapter 17). Satiety and nausea are, in this sense, neighbouring outputs of the same accessible brainstem region, which is why they often travel together early in treatment and why the appetite benefit and the gastrointestinal side effect share a mechanistic root.

6.4 "Food noise" — the clinical phenomenon

One of the most consistent and striking things patients report is not simply reduced hunger but the disappearance of what has come to be called **food noise** — the constant, intrusive, background preoccupation with food, the mental chatter about what and when to eat next. Patients describe a quietening of this preoccupation that they often find more transformative than the reduction in hunger per se. Mechanistically, food noise sits at the interface of the homeostatic and hedonic systems: it reflects both a reduction in the drive to eat (this chapter) and a reduction in the reward salience of food (Chapter 7). Its disappearance is part of why patients on semaglutide reduce intake without the sense of effortful restriction that characterises dieting — the constant pull toward food is simply diminished. This subjective transformation is clinically valuable because it is sustainable in a way that willpower-based restriction is not, and it is a useful thing to ask patients about directly, as its presence often correlates with response.

07 MECHANISM V

Reward, Cravings and Addiction: Food, Alcohol and Nicotine

One of the most clinically striking effects patients describe is not the reduction of hunger but the quieting of **craving** — the disappearance of intrusive, magnetic thoughts about food, and, in a rapidly growing literature, a reduced desire for alcohol and cigarettes. This is a mechanism distinct from gastric emptying or hypothalamic satiety, and it explains effects that the homeostatic model of Chapter 6 cannot. The mechanism is the action of GLP-1-receptor signalling within the brain's **reward circuitry**, and it has become one of the most scientifically exciting and clinically consequential aspects of the drug class — extending its potential reach from metabolic disease into the territory of substance-use disorders. This chapter develops the reward circuitry in depth, explains how semaglutide reaches and dampens it, and reviews the current evidence across food, alcohol and nicotine, with appropriate care about what is established and what remains investigational.

7.1 Why craving is a distinct mechanism

It is essential to separate two things that everyday language conflates: *hunger* and *craving*. Hunger is the homeostatic drive to eat that arises from energy need and is governed by the hypothalamic and gut mechanisms of Chapter 6. Craving is the hedonic pull toward a specific rewarding substance or experience — a palatable food, a drink, a cigarette — that arises largely independent of energy need and is governed by the brain's reward system. A person can crave dessert when entirely full; that is hedonic, not homeostatic. The reason this distinction matters pharmacologically is that semaglutide reduces both, through different circuitry, and the craving effect in particular cannot be explained by slowed gastric emptying or by POMC/AgRP satiety signalling. Patients who report that food "lost its grip," that they can take or leave a drink they previously could not resist, or that cigarettes lost their appeal, are describing a change in the reward system, and that is the subject of this chapter.

7.2 The mesolimbic reward circuit

The anatomical seat of reward and craving is the **mesolimbic dopamine system**. Its origin is the **ventral tegmental area** (VTA) of the midbrain, which contains dopamine-producing neurons that project to two principal targets: the **nucleus accumbens** (NAc) in the ventral striatum, and the **prefrontal cortex** (PFC). When a reward — or, crucially, a *cue* that predicts a reward — is encountered, VTA dopamine neurons fire and release dopamine in the nucleus accumbens, and this dopamine signal assigns **incentive salience** to the cue: it makes the cue "wanted," it grabs attention and motivates approach and consumption. Modern neuroscience distinguishes this dopamine-mediated "*wanting*" (incentive salience, the motivational pull) from the more opioid-mediated "*liking*" (the hedonic pleasure of consumption itself); addiction and compulsive consumption are characterised especially by excessive *wanting* — intense, cue-triggered craving — often out of proportion to the actual pleasure obtained. The dopamine signal also encodes **reward prediction**, learning to fire to the cues that predict reward, which is how environmental triggers (the sight of food, the bar, the pack of cigarettes) acquire their power to provoke craving. This circuit is the common final pathway for the reinforcing

properties of food, alcohol, nicotine and other drugs of abuse, which is why a single intervention acting on it can influence several behaviours at once.

7.3 GLP-1 receptors sit inside the reward system

The key enabling fact is that GLP-1 receptors are not confined to the pancreas and hypothalamus but are expressed *throughout* the mesolimbic and mesocortical reward circuitry. They are found on or near neurons of the **ventral tegmental area**, the **nucleus accumbens**, the **prefrontal cortex**, the lateral septum, and the hindbrain **nucleus of the solitary tract** (NTS). Moreover, as introduced in Chapter 1, the brain contains its own population of GLP-1-producing neurons — the preproglucagon neurons of the NTS — which project to precisely these reward regions, meaning the reward system is natively wired to be modulated by GLP-1 signalling. The drug is therefore not acting on an incidental, off-target site; it is engaging a physiological GLP-1 system whose normal role includes the regulation of the reward value of ingested substances. This anatomical fact is the foundation of everything that follows.

7.4 How semaglutide reaches and dampens the circuit

A long-acting, peripherally administered agonist reaches these deep brain regions by two routes that bypass the need to cross an intact blood–brain barrier (the same routes described in Chapter 6): **vagal afferents** synapsing in the NTS, and direct access at **circumventricular organs** such as the area postrema and median eminence, where the barrier is incomplete. Once engaged, GLP-1-receptor activation *inhibits the dopamine neurons of the VTA* and reduces stimulated dopamine release in the nucleus accumbens. The experimental signature, demonstrated across numerous preclinical studies, is that the dopamine surge normally evoked in the accumbens by food, alcohol, nicotine and even cocaine is *blunted* by GLP-1-receptor agonism. The functional consequence is elegant: the rewarding cue still occurs, but it now predicts and produces a smaller dopamine signal, so the incentive salience — the "wanting" — attached to it shrinks. The pull of the cue weakens.

Semaglutide also acts on this system by opposing **ghrelin**, the stomach-derived "hunger hormone" that not only stimulates appetite but amplifies mesolimbic reward signalling, increasing the incentive value of food and other rewards via its own receptors in the VTA. By counteracting ghrelin's pro-reward input while adding its own anti-reward (dopamine-dampening) signal, semaglutide shifts the balance of the reward system away from craving from two directions at once. More broadly, central GLP-1-receptor activation modulates not only dopaminergic but also glutamatergic and GABAergic signalling within these circuits, reshaping the reward and motivational response to reinforcing cues. The net result, at the level of behaviour, is reduced craving and reduced consumption of reinforcing substances.

PEARL Patients describe this as the end of "food noise." The same circuit-level change — reduced incentive salience of a reinforcing cue — is the most parsimonious explanation for the parallel reductions seen in alcohol and tobacco use. It is one mechanism with several behavioural read-outs, which is why a metabolic drug is being investigated across multiple addictions.

08 PHARMACOKINETICS

Absorption, Distribution, Metabolism and Elimination

With the mechanisms established, the monograph turns to the practical pharmacology that governs how semaglutide is dosed and used. Pharmacokinetics — what the body does to the drug — is not an abstract topic here but 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 why the oral formulation demands such a strict administration ritual. 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.

Parameter	Subcutaneous	Oral (with SNAC)
Half-life	≈ 7 days	≈ 7 days
Time to steady state	4–5 weeks	4–5 weeks
Bioavailability	≈ 89%	≈ 0.4–1% (highly variable)
Time to peak (T_{max})	1–3 days	~1 hour (absorption window)
Plasma protein binding	> 99% (reversibly albumin-bound)	
Dosing interval	Once weekly	Once daily
Metabolism	Proteolysis + fatty-acid β-oxidation; no major CYP route	
Renal/hepatic impairment	No dose adjustment required; limited data in severe impairment	

8.1 Absorption — subcutaneous and oral

After subcutaneous injection, semaglutide is absorbed slowly and almost completely, with an absolute bioavailability of approximately 89% and a time to peak plasma concentration (T_{max}) of one to three days. The slowness of absorption is itself a consequence of the albumin-binding design: the drug is held in a depot at the injection site and in the circulation rather than appearing as a sharp peak. This contributes to the smooth, low-variability concentration profile that is one of the molecule's clinical virtues, since the absence of sharp peaks reduces concentration-related side effects.

Oral absorption is an entirely different story, governed by the SNAC co-formulation described in Chapter 2. Absolute oral bioavailability is very low — on the order of 0.4–1% — and, crucially, variable, because it depends so sensitively on the conditions of administration. Absorption occurs principally in the *stomach*, within the brief, localised window that SNAC creates by raising local pH (protecting the peptide and inactivating pepsin) and promoting transcellular gastric uptake. The oral T_{max} is short (around an hour), reflecting this gastric absorption window. Because the window is so sensitive to disruption, food, more than a small

sip of water, and co-ingested substances markedly reduce absorption — the mechanistic basis of the strict empty-stomach administration rules of Chapter 10. Despite these differences in *how* the drug is absorbed, once it reaches the systemic circulation the oral and subcutaneous forms behave identically, because they are the same molecule.

8.2 Distribution and the albumin reservoir

Once absorbed, semaglutide is more than 99% bound to plasma **albumin**, reversibly, through the fatty-diacid side chain (Chapter 2). The volume of distribution is correspondingly small — on the order of 8–12 litres, reflecting that the drug is largely confined to the plasma and extracellular fluid rather than being widely distributed into tissues. This albumin binding is the central pharmacokinetic fact of the molecule: it is 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 is what diffuses to target tissues — including, by the routes described in Chapters 6 and 7, the relevant brain regions — to engage the GLP-1 receptor.

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

The elimination half-life of semaglutide is approximately **one week** (around 7 days) for both the subcutaneous and oral routes. Three clinically important consequences follow directly from this single number. First, because steady state is reached after roughly four to five half-lives, plasma concentrations continue to rise for about **4–5 weeks** after starting or after each dose increase. This is why the effect — and the tolerability — at any given dose cannot be fully judged for several weeks, and why titration is deliberately gradual (Chapter 10): the drug is still accumulating toward its steady-state level. Second, the week-long half-life buffers the consequences of a **missed dose**: because concentrations decline slowly, a missed weekly subcutaneous dose can be taken within about five days of the scheduled day without materially disrupting therapy, after which the patient resumes the regular schedule. Third, and conversely, the same slow decline means that effects — including adverse effects, and including any pharmacological action relevant to pregnancy — **persist for several weeks after the last dose**. There is no rapid washout, which is why discontinuation before a planned pregnancy or before certain procedures must be timed well in advance rather than assumed to clear quickly.

> **PEARL** Almost every dosing rule for semaglutide is a restatement of its ~1-week half-life. Slow accumulation over 4–5 weeks → gradual titration and delayed judgement of effect. Slow decline → a late dose is tolerated, but the drug lingers for weeks after stopping (relevant to pregnancy planning and the peri-operative period).

8.4 Metabolism — why there are almost no CYP interactions

Semaglutide is a peptide, and it is metabolised as a peptide rather than by the cytochrome P450 (CYP) system that handles most small-molecule drugs. Elimination proceeds by two routes intrinsic to the molecule's structure: *proteolytic cleavage* of the peptide backbone by ubiquitous peptidases (including, where accessible, neprilysin and DPP-4-type activities), and *β-oxidation* of the fatty-acid side chain, the same metabolic process the body uses to break

09 SELECTION

Patient Selection — For Whom, and For Whom Not

The mechanisms of Chapters 3 to 7 dictate the ideal candidate for semaglutide: someone whose disease is driven by insulin resistance, excess adiposity, or the atherosclerotic and renal risk that sits on that background. The licensed indications map onto exactly those mechanisms, which is why a mechanistic understanding makes patient selection almost self-evident rather than a matter of memorised criteria. This chapter sets out who the drug is for, who benefits most, who should not receive it and why, the special situations that demand caution, and the pre-prescription assessment that should precede the first dose. It closes the gap between the pharmacology and the clinic: the right patient is the one whose pathophysiology matches the drug's mechanism and whose risk of the feared adverse events (Chapter 17) is acceptable.

9.1 The licensed indications

Semaglutide's indications have expanded rapidly from glycaemic control to a series of hard-outcome and organ-protective uses, each corresponding to a mechanism developed earlier (Table 9.1). For **type 2 diabetes**, semaglutide is an adjunct to diet and exercise, available subcutaneously and orally, drawing on the insulin-secretory, glucagon-suppressing and weight-loss mechanisms of Chapters 3 to 6. For **chronic weight management**, the high-dose regimen 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 appetite and reward mechanisms (Chapters 6 and 7) and the STEP trials. For **cardiovascular risk reduction**, semaglutide is indicated in people with established cardiovascular disease and overweight or obesity but *without* diabetes, following SELECT (Chapter 13) — an indication that reframed obesity itself as a modifiable cardiovascular risk factor. For **chronic kidney disease in type 2 diabetes**, the FLOW trial (Chapter 14) established a renal-protective indication. And in **metabolic dysfunction-associated steatohepatitis (MASH)**, the ESSENCE programme (Chapter 15) supports an emerging hepatic indication. The breadth of these indications is not a series of lucky discoveries but the predictable clinical harvest of correcting the integrated cardiometabolic phenotype (Chapter 6).

Table 9.1. Approved indications and thresholds (semaglutide).

Indication	Threshold / population	Evidence
Type 2 diabetes	Adjunct to diet/exercise; SC or oral	SUSTAIN, PIONEER[10,11]
Chronic weight management	BMI ≥ 30, or ≥ 27 with ≥ 1 weight-related comorbidity	STEP programme[6,7]
CV risk reduction	Established CVD with overweight/obesity, *without* diabetes	SELECT[12]
Chronic kidney disease in T2D	T2D with CKD	FLOW[13]
MASH (emerging)	Biopsy-confirmed steatohepatitis with fibrosis	ESSENCE[15]

9.2 Who benefits most

Within and beyond the formal indications, the patients who benefit most are those whose clinical picture is dominated by the mechanisms the drug addresses. The patient with **type 2 diabetes and obesity** gains glycaemic control, weight loss and cardiometabolic improvement simultaneously. The patient with **established atherosclerotic disease and excess weight** gains cardiovascular event reduction. The patient with **diabetic kidney disease** gains slowing of renal decline. The patient with **obesity and a craving-driven or hedonic eating pattern** (Chapter 7) gains a reduction in the food reward that drives their intake, often with the added benefit of reduced hazardous drinking or smoking. In general, benefit is largest where insulin resistance and adiposity dominate the picture, and the prescriber should think in terms of matching the patient's predominant mechanism to the drug rather than reaching for it reflexively. Conversely, where hyperglycaemia is driven overwhelmingly by beta-cell failure rather than insulin resistance (advanced, lean type 2 diabetes), the insulin-secretory contribution is smaller (Chapter 3) and expectations should be adjusted.

9.3 Contraindications and major cautions

The situations in which semaglutide should be avoided or used only with caution (Table 9.2) follow from its mechanism and its adverse-effect profile (Chapter 17). The single formal class **contraindication** is a personal or family history of **medullary thyroid carcinoma** (MTC) or the **multiple endocrine neoplasia type 2** (MEN-2) syndrome. This derives from rodent studies in which GLP-1-receptor agonists caused thyroid C-cell tumours; although large human population analyses have been reassuring and have not confirmed a meaningful human MTC signal (Chapter 17), the boxed contraindication remains and should be respected, because the consequence of a missed inherited C-cell tumour syndrome is severe and the contraindication costs little to honour.

Several **cautions**, rather than absolute contraindications, require judgement. A history of **pancreatitis** warrants caution and a weighing of risk (the signal is real but modest and partly confounded, Chapter 17). **Severe gastroparesis** or a significant gut-motility disorder is a caution because the drug's delayed gastric emptying would compound the problem. **Active proliferative diabetic retinopathy** warrants co-management with ophthalmology and attention to the rate of glucose lowering, because rapid improvement in glycaemia can transiently worsen retinopathy (the SUSTAIN-6 early-worsening signal, Chapter 16). A prior

10 DOSING I

Titration — Subcutaneous and Oral

The governing principle of semaglutide dosing is **slow titration**: gastrointestinal tolerability, not efficacy, sets the pace. Almost every adverse event occurs during dose escalation and almost every one resolves with time at a stable dose. A prescriber who understands *why* titration is necessary — and it is a direct consequence of the mechanisms of Chapters 6 and 8 — will titrate confidently, neither rushing escalation (which causes avoidable nausea and dropouts) nor stalling unnecessarily at a sub-therapeutic dose. This chapter sets out the escalation schedules for each formulation, the administration technique, how to manage the titration phase, what to do about missed doses and switching, and the co-medication adjustments that must accompany initiation.

10.1 Why titration is mandatory

The need to start low and go slow rests on two mechanistic facts. The first is that the dose-limiting adverse effects — nausea, vomiting, and the gastrointestinal symptoms of Chapter 17 — arise from the same actions that produce benefit: engagement of the area postrema (the accessible hindbrain emetic and satiety sensor, Chapter 6) and slowing of gastric emptying. These effects are *dose-related*, so introducing a high dose abruptly provokes intense symptoms. The second fact is that the relevant effects show **tachyphylaxis** — they attenuate with continued exposure as the patient adapts. Gradual escalation exploits this adaptation: each small step is tolerated, the patient adapts over the weeks at that step, and the next step is then tolerated in turn. Layered on top is the pharmacokinetic reality (Chapter 8) that, with a one-week half-life, each dose takes about four to five weeks to reach steady state, so a roughly four-week interval per step both allows adaptation and lets each dose approach its full effect before the next increase. The 0.25 mg starting dose is therefore an *initiation (tolerability) dose only*, not a therapeutic one — its purpose is to begin the adaptation, not to treat.

10.2 Subcutaneous semaglutide — weight management

For chronic weight management, the standard schedule escalates at roughly four-week intervals: **0.25 → 0.5 → 1.0 → 1.7 → 2.4 mg** once weekly, with 2.4 mg the maintenance/target dose. The 0.25 mg dose is initiation only. The cardinal rule of titration is to *escalate only when the current dose is adequately tolerated*: if symptoms are marked at a given step, hold at that step (or step back down) for an extra few weeks and attempt the increase again once symptoms settle, rather than pushing on. The four-week interval is a minimum, not a mandate — it is entirely acceptable, and often wise, to spend longer at a step. A patient who reaches a tolerated dose below the maximum that delivers adequate clinical benefit need not always be pushed to the ceiling; the target is the lowest dose achieving the goal with acceptable tolerability.

10.3 Subcutaneous semaglutide — type 2 diabetes

For type 2 diabetes, the schedule escalates **0.25 → 0.5 → 1.0 → 2.0 mg** weekly, again in approximately four-week steps, titrated to the glycaemic target and to tolerability. As in

weight management, escalation is governed by tolerance, and the dose is advanced only as far as is needed to reach the glycaemic goal. Many patients achieve good control at 0.5 or 1.0 mg without needing the maximum. Where weight loss is also a goal in a patient with diabetes, the higher doses used for weight management may be appropriate within the relevant product's licensing.

10.4 Oral semaglutide and the empty-stomach ritual

Oral semaglutide (for type 2 diabetes) escalates **3 → 7 → 14 mg** once daily in roughly 30-day steps, the 3 mg dose being for initiation only. Higher-strength oral formulations for weight management entered use in 2025 (Chapter 19). For the oral formulation, absorption rather than tolerability is the limiting factor, and the administration rule is absolute, not advisory: take the tablet on an **empty stomach on waking**, with **no more than about 120 mL of plain water**, and wait **at least 30 minutes** before any food, other fluid, or other oral medication. Swallow the tablet whole; do not split, crush or chew it.

> **The empty-stomach rule is mechanistic, not optional.** SNAC raises the local gastric pH to protect the peptide and inactivate pepsin, creating the brief gastric absorption window (Chapter 2); food, excess water or co-administered drugs dilute and disperse this effect and can abolish absorption. A patient who takes oral semaglutide with breakfast is effectively untreated. Because the absorbed fraction is small (~1%) and variable, *consistency* of correct administration is as important as the dose itself.

10.5 Administration technique

The subcutaneous formulation is given by once-weekly injection into the abdomen, thigh or upper arm; injection sites should be rotated to reduce local reactions, and the same weekday should be kept for consistency (the day may be changed when needed provided the interval between doses is at least about 48–72 hours). The pre-filled pen requires appropriate cold-chain storage before first use, after which limited time at room temperature is permitted per the product instructions; the peptide is sensitive to heat and to freezing, and a pen that has been frozen or overheated should not be used (Chapter 2). For the oral tablet, the ritual of Section 10.4 *is* the technique, and teaching it carefully — ideally with a teach-back from the patient — is one of the highest-yield things a prescriber can do to ensure the oral drug actually works.

10.6 Managing the titration phase

The titration phase is where the therapeutic relationship is won or lost, because it is when symptoms are worst and adherence most fragile. The pace should be individualised: hold or step down rather than forcing escalation in the face of significant symptoms, and reassure the patient 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 and are developed in Chapters 17 and 18. Adequate hydration is important, particularly if vomiting or diarrhoea occurs, both for comfort and to protect renal function (Chapter 17). Setting expectations in advance — that

11 DOSING II

Minimum Effective Dose, Plateaus, and the Current Landscape

Chapter 10 set out how to escalate semaglutide; this chapter addresses the harder, more individual questions that follow once a patient is established on therapy: what is the right dose for *this* patient and goal, how low can a maintenance dose be, why does weight loss stall, and how should the dose be used to break a plateau. It then defines the off-label practices of micro- and macro-dosing, and sets out the mid-2026 regulatory and supply landscape that shapes how dosing actually occurs in practice. The unifying theme is that the dose should be chosen by *clinical reasoning* — by goal, response and tolerability — rather than reflexively driven to the maximum or reflexively minimised.

11.1 The minimum effective dose principle

The single most useful dosing principle, and one frequently misunderstood in both directions, is that the goal of dosing is the **minimum effective dose** — the lowest dose that achieves the patient's clinical objective with acceptable tolerability — not the maximum dose the label permits. The titration schedule of Chapter 10 climbs toward a ceiling, but the ceiling is a *limit*, not a *target*. Each dose level exists for a reason, and the dose a given patient should be on is determined by what that patient needs, not by how high the drug can go.

This principle has a firm mechanistic and clinical basis. The dose–response relationship is real but not linear in benefit-versus-harm: higher doses produce greater average weight loss and glycaemic effect, but they also produce more gastrointestinal toxicity, and the incremental benefit of each step diminishes while the incremental burden does not. Consequently, many patients reach their glycaemic or weight goal well below the maximum dose, and for them there is no clinical reason — and a real tolerability cost — to climb further. The corollary, sometimes called *reason dosing*, is that every dose decision should have an articulable clinical reason: "escalate because the goal is not yet met and the current dose is tolerated," "hold because the goal is met," "hold because tolerability is marginal," or "this is the lowest dose that maintains the result." A dose chosen by reflex — always to the maximum, or always to the minimum — is a dose chosen without reasoning.

For *maintenance* specifically, once a patient has reached and stabilised at their goal, the aim becomes the lowest dose that *holds* the result with good tolerability, which may be at or below the dose that achieved it. An important honesty here: the registration trials studied fixed target doses, and the evidence for deliberately down-titrating to a lower maintenance dose after success is limited, with a recognised risk that reducing the dose allows partial weight regain as the defended set-point reasserts itself (Chapter 6 and Chapter 18). The defensible practice is therefore to individualise — to use the lowest dose that demonstrably maintains the goal for that patient, monitoring for regain — rather than to assume that everyone must remain indefinitely at the maximum, or that everyone can safely be reduced.

PEARL — reason dosing
The labelled maximum is a ceiling, not a target. Ask of every dose: *what is the clinical reason this patient is on this dose?* If the goal is met and tolerated, there is no reason to climb. If the goal is unmet and the dose is tolerated and below maximum, there is a reason to climb. The plateau question below is simply this principle applied to a stall.

11.2 The weight-loss plateau — why it happens

Almost every patient on semaglutide eventually reaches a **weight-loss plateau** — a point at which weight stabilises and further loss stops. Understanding why is essential to responding correctly, because the cause determines whether a dosing change is the right answer. The plateau is, in most cases, a predictable physiological equilibrium rather than a failure. As weight falls, two homeostatic responses (Chapter 6) progressively oppose further loss: total energy expenditure declines (both because a smaller body burns fewer calories and because of **adaptive thermogenesis**, a disproportionate fall in expenditure), and the body's defended set-point partially reasserts itself, so appetite and the drive to eat recover somewhat from their initially suppressed state. As intake and expenditure re-equilibrate at the new, lower weight, loss ceases and a plateau is reached. At a fixed drug dose, this new equilibrium is the expected endpoint, not a sign that the drug has "stopped working."

Before treating a plateau as a true equilibrium, however, two *pseudo-plateaus* must be excluded, because they call for patience rather than action. First, an apparent stall during ongoing titration, or in the weeks just after a dose increase, may simply reflect that the current dose has *not yet reached steady state* — recall that each dose takes about four to five weeks to approach its full effect (Chapter 8), so the benefit of a recent increase may not yet be fully expressed. Second, an apparent stall in a patient still on a sub-maximal dose is not a true ceiling at all, 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 two pseudo-plateaus is the first analytical step, and it leads directly to the dosing strategy below.

11.3 Breaking a plateau with dosing

Once a genuine plateau is identified, the dosing approach follows from the minimum-effective-dose principle and the dose–response relationship, and it divides cleanly according to whether the patient is already at the maximum dose.

If the patient is below the maximum dose and tolerating it, the primary dosing lever is *escalation to the next step*. The registration programmes demonstrate a genuine dose–response for weight loss — higher maintenance doses produce greater average loss — so a patient who has plateaued at, say, 1.0 or 1.7 mg, and who is tolerating that dose, 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 and best-evidenced way to break a plateau: the stall is, in effect, the signal that the current dose has delivered what it can, and the dose-response headroom remains. Escalation should still follow the tolerability rules of Chapter 10 — advance only if the current dose is tolerated, and allow the new dose four to five weeks to express its effect before judging it.

12 EFFICACY I

Glycaemic and Weight Efficacy

The mechanisms of the preceding chapters predict that semaglutide should lower glucose substantially and reduce weight to a degree no previous medical therapy had achieved. The trial programmes confirm both. This chapter reviews the weight-loss evidence (the STEP programme), the glycaemic evidence (SUSTAIN for the subcutaneous formulation, PIONEER for the oral), the magnitude of these effects in context, the composition and durability of the weight loss, and — with appropriate honesty — the head-to-head comparisons in which the dual agonist tirzepatide proves superior. The figures place the numbers in visual context.

12.1 Weight — the STEP programme

The **STEP** (Semaglutide Treatment Effect in People with obesity) programme established semaglutide 2.4 mg as the first medical therapy to produce weight loss approaching the range previously achievable only with surgery. In the pivotal **STEP-1** trial, adults with obesity but without diabetes lost approximately **14.9%** of body weight over 68 weeks on semaglutide 2.4 mg, versus about 2.4% on placebo (both with lifestyle support) — a placebo-subtracted difference of roughly 12 percentage points, and a degree of pharmacological weight loss without precedent at the time. The loss was progressive over the titration and maintenance period rather than abrupt, as the dose climbed and appetite suppression took hold (Figure 12.1).

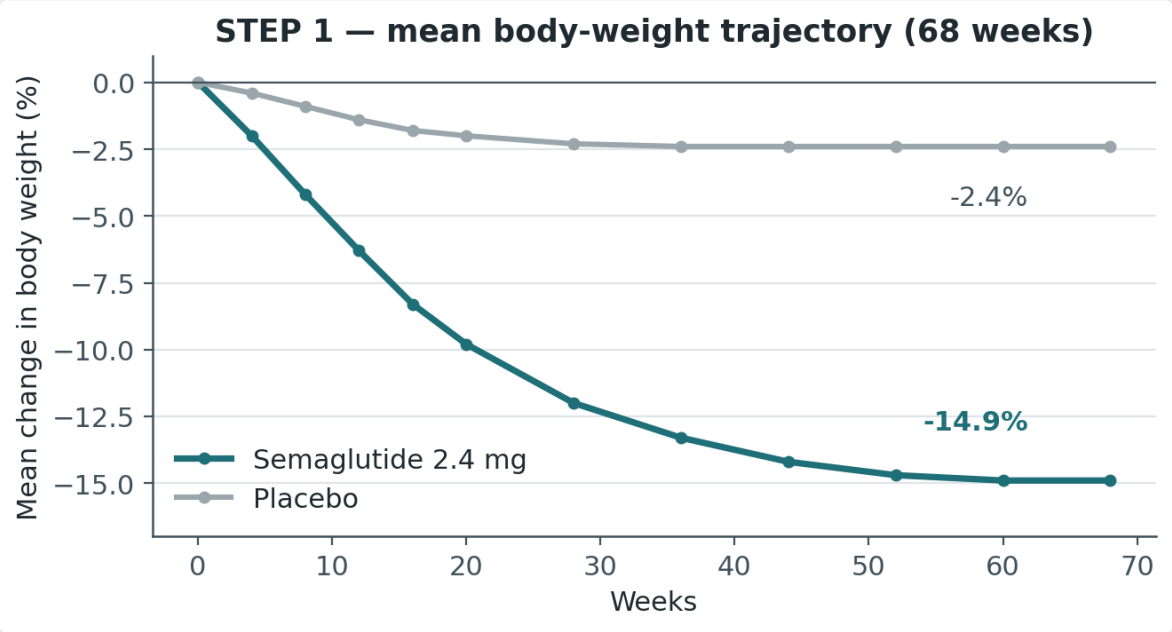


Figure 12.1. STEP-1: progressive weight loss with semaglutide 2.4 mg versus placebo over 68 weeks.

The wider programme refined and reinforced this result. **STEP-2** studied patients with type 2 diabetes, in whom weight loss is characteristically somewhat smaller (around 10%), reflecting the more complex metabolic milieu. **STEP-3** added intensive behavioural therapy, achieving somewhat greater loss and underscoring that the drug and lifestyle are complementary (Chapter 18). **STEP-8** compared semaglutide head-to-head with the earlier GLP-1 agonist liraglutide and demonstrated clearly superior weight loss. And the durability trials, **STEP-4**

and **STEP-5**, are discussed in Section 12.5 because they carry the most important practical lesson of the programme. Across these trials, a consistent and clinically meaningful proportion of patients achieved 15%, and many achieved 20%, weight loss — thresholds at which obesity-related comorbidities materially improve.

12.2 Magnitude in context

To appreciate what a ~15% mean weight loss means, it must be placed against the alternatives (Figure 12.2). Sustained lifestyle intervention alone typically produces only a few percent weight loss on average, much of it regained. The earlier generation of weight-loss pharmacotherapy — including liraglutide — produced mean losses in the region of 5–8%. Semaglutide 2.4 mg roughly doubled that, into the 15% range, while the dual agonist tirzepatide and the investigational triple agonists extend it further still (toward 20% and beyond). Bariatric surgery remains more powerful again (often 25–30%), so semaglutide does not abolish the role of surgery, but it has dramatically narrowed the gap and created an effective middle tier of therapy that did not previously exist. The clinical significance is not merely cosmetic: at these magnitudes, the integrated cardiometabolic phenotype (Chapter 6) improves across the board, which is what the hard-outcome trials of Chapters 13 to 15 then demonstrate.

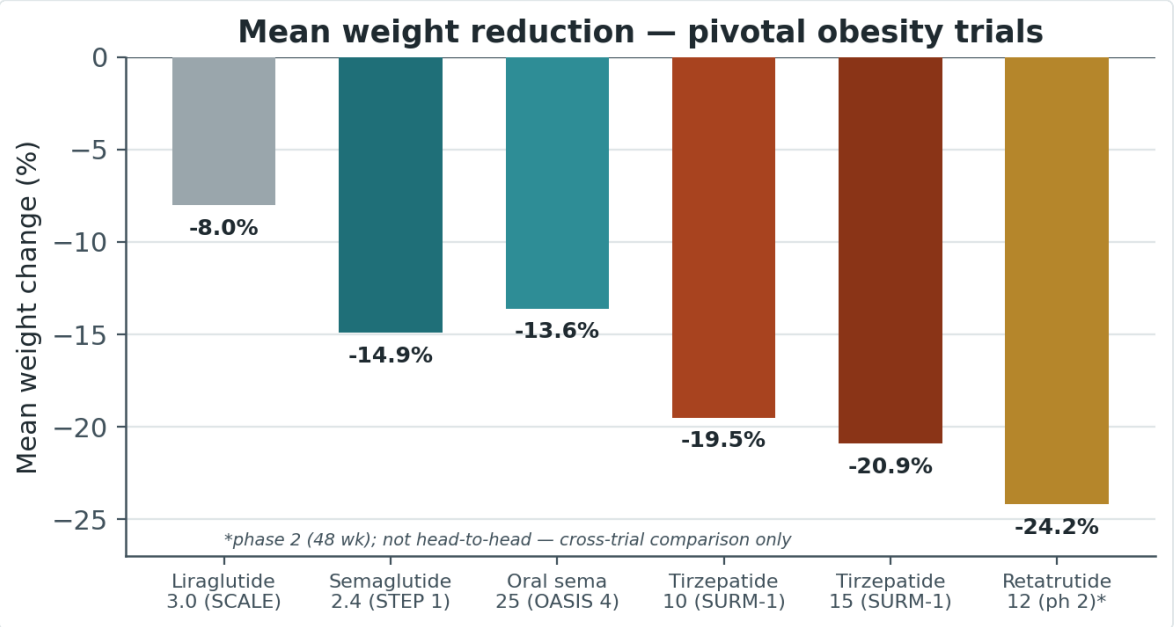


Figure 12.2. Approximate mean weight loss across agents in obesity trials — placing semaglutide between earlier GLP-1 agonists and the dual/triple agonists.

12.3 Glycaemia — SUSTAIN and PIONEER

For type 2 diabetes, two programmes established semaglutide's glycaemic efficacy. The **SUSTAIN** programme tested the subcutaneous formulation and showed robust HbA1c reductions — on the order of 1.5–1.8 percentage points at the higher doses, among the largest of any glucose-lowering class — together with weight loss, reflecting the combined insulin-secretory, glucagon-suppressing and weight mechanisms of Chapters 3 to 6. Importantly, in the head-to-head **SUSTAIN-7** trial, semaglutide produced greater reductions in both HbA1c and weight than the weekly GLP-1 agonist dulaglutide, establishing it at the top of the

13 OUTCOMES I

Cardiovascular Outcomes and Mechanism

Glycaemic and weight efficacy (Chapter 12) matter because of what they prevent. The cardiovascular outcome trials are where the integrated phenotype of Chapter 6 translates into the endpoints that matter most to patients — heart attacks, strokes and deaths — and where semaglutide moved from a metabolic drug to a cardioprotective one. This chapter reviews those trials, examines the striking finding that much of the benefit appears *before* and *independent of* major weight loss, and develops the mechanisms that explain it. The central and clinically liberating conclusion is that the cardiovascular protection is a property of GLP-1-receptor signalling itself, not merely a by-product of being lighter.

13.1 The outcome trials

In type 2 diabetes, **SUSTAIN-6** reduced three-point major adverse cardiovascular events (MACE — cardiovascular death, non-fatal myocardial infarction, non-fatal stroke) with a hazard ratio of 0.74, and the oral **PIONEER-6** trial was consistent (HR 0.79). These established cardiovascular safety and a benefit signal in the diabetic population. The landmark, however, is **SELECT**: in patients with established cardiovascular disease and overweight or obesity but **without diabetes**, semaglutide 2.4 mg reduced MACE by **20% (HR 0.80)** over a mean of several years. This result was pivotal not only for its magnitude but for its population — it demonstrated cardiovascular benefit in people whose indication was obesity, not diabetes, thereby establishing *obesity itself as a modifiable cardiovascular risk factor* and extending the drug's indication into cardiovascular prevention. Figure 13.1 places these hazard ratios alongside the renal outcomes of Chapter 14.

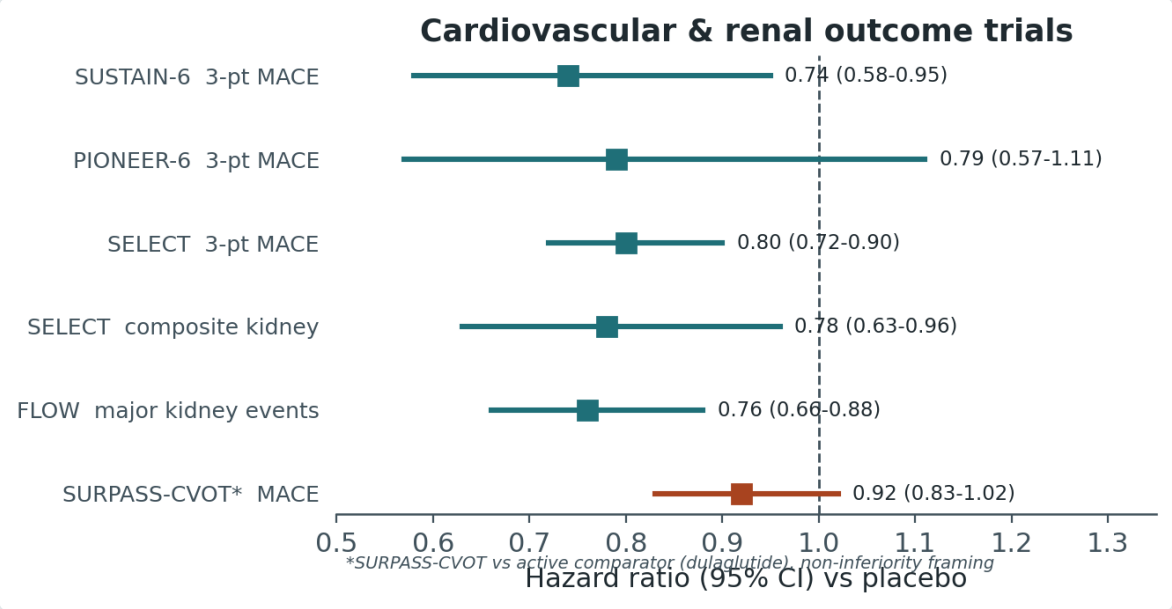


Figure 13.1. Cardiovascular and renal outcome trials for semaglutide (with SURPASS-CVOT shown for context). Hazard ratios versus placebo; SURPASS-CVOT is versus active comparator.

13.2 SELECT in detail

SELECT deserves closer attention because of its conceptual importance. The trial enrolled patients with prior cardiovascular disease and a BMI of 27 or above but without diabetes, and followed them for cardiovascular events. Beyond the 20% MACE reduction, several features are notable. The benefit was accompanied by improvements across the cardiometabolic profile — blood pressure, lipids, inflammatory markers and waist circumference — consistent with correction of the integrated phenotype. Among the large prediabetic subgroup, semaglutide produced a striking **73% relative reduction in progression to type 2 diabetes**, a powerful demonstration of disease prevention that has its own long-term cardiovascular and microvascular implications. And the consistency of benefit across subgroups reinforced that the effect was a genuine property of the therapy rather than confined to a particular patient type. SELECT thus did more than earn an indication; it reframed how the cardiovascular risk of obesity should be understood and treated.

13.3 Why — the mechanisms of cardiovascular benefit

The most scientifically interesting and clinically consequential aspect of the cardiovascular data is *how* the benefit arises. Two observations point away from weight loss as the sole explanation. First, the benefit emerged **early**, before much weight had been lost — the event curves began to separate sooner than a weight-mediated mechanism alone would predict. Second, prespecified mediation analyses showed the effect was largely **independent of baseline HbA1c and of the magnitude of weight loss**: change in waist circumference statistically mediated only about a third of the effect, leaving the majority attributable to other, more direct actions of GLP-1-receptor signalling. Several mechanisms operate together (Table 13.1).

Table 13.1. Mechanisms of cardiovascular benefit.

Pathway	Effect
Anti-inflammatory	hsCRP fell ~38% in SELECT; in animal models semaglutide reverses atherosclerotic gene programmes (leukocyte adhesion, plaque) — partly weight/cholesterol-independent[17,18]
Endothelial / vascular	Improved endothelial function, vasodilation, reduced vessel-wall remodelling
Plaque stability	Reduced lesion progression and inflammatory infiltration
Haemodynamic / metabolic	Lower blood pressure, improved lipids, reduced epicardial/visceral fat
Glycaemic (diabetes prevention)	In SELECT, 73% relative reduction in progression to diabetes among prediabetics

The **anti-inflammatory** arm is probably central. Atherosclerosis is fundamentally an inflammatory disease of the vessel wall, and GLP-1 receptors are expressed on the immune and vascular cells involved. The ~38% fall in high-sensitivity CRP in SELECT is a marker of a genuine reduction in vascular inflammation, and preclinical work (for example, the demonstration that semaglutide reverses atherosclerotic gene-expression programmes and reduces plaque in animal models, partly independently of weight and cholesterol) supports a direct anti-atherosclerotic action. Layered on this are **endothelial** improvements (better

14 OUTCOMES II

Renal Outcomes and Their Mechanism — FLOW

Diabetic kidney disease is the leading cause of kidney failure worldwide and a powerful multiplier of cardiovascular risk, yet for decades the only disease-modifying therapies were blockade of the renin–angiotensin system and, more recently, the SGLT2 inhibitors and finerenone. The demonstration that semaglutide slows the progression of kidney disease adds a further pillar to this framework, and it was long suspected from the consistent reductions in albuminuria seen across the trial programme before being confirmed by a dedicated outcome trial. This chapter reviews that evidence and develops the mechanisms of renal protection, which — as with the cardiovascular benefit — are only partly explained by weight loss.

14.1 The evidence

The **FLOW** trial randomised patients with type 2 diabetes and chronic kidney disease to semaglutide or placebo and was *stopped early for efficacy* — a strong signal of benefit. It showed a **24% reduction** in major kidney events (a composite of kidney failure, a substantial and sustained loss of eGFR, and renal or cardiovascular death), with a hazard ratio of **0.76**. The benefit spanned both the renal-specific components and the cardiovascular death component, underlining the tight coupling of kidney and cardiovascular disease in this population. Consistent with FLOW, the kidney composite endpoint in **SELECT** pointed the same way (hazard ratio around 0.78), and earlier trials had shown reductions in albuminuria — an established marker of both renal and cardiovascular risk — across the programme. These converging findings (see also Figure 13.1) establish a genuine, disease-modifying renal benefit rather than a mere metabolic epiphenomenon.

14.2 Why — mechanisms of renal protection

Several mechanisms, again only partly weight-mediated, converge to protect the kidney (Table 14.1), and understanding them clarifies why the benefit is real and how it complements other renoprotective agents.

Table 14.1. Mechanisms of renal benefit.

Pathway	Effect
Reduced glomerular hyperfiltration	Lower intraglomerular pressure → less mechanical podocyte/capillary 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, weight and improved lipids reduce the drivers of progression

The **haemodynamic** arm is probably central and is worth spelling out. In diabetes and obesity, the glomerulus is subjected to **hyperfiltration** — an elevated intraglomerular

pressure that, over time, mechanically injures the podocytes and capillary wall, driving albuminuria and progressive scarring. GLP-1-receptor signalling reduces this hyperfiltration, partly through a natriuretic action: by reducing sodium reabsorption in the proximal tubule, it increases sodium delivery to the distal nephron, which — through tubuloglomerular feedback — signals the afferent arteriole to constrict and lowers intraglomerular pressure. This is mechanistically analogous to (though distinct in detail from) the way SGLT2 inhibitors protect the kidney, and it is part of why the two classes are complementary. The natriuresis also contributes to the blood-pressure reduction that itself protects the kidney. Layered on the haemodynamic effect are direct **anti-inflammatory and anti-fibrotic** actions in the kidney (reducing the inflammatory and fibrotic signalling that drives progression), and the broad **systemic risk-factor control** — lower glucose, blood pressure, weight and improved lipids — that removes the upstream drivers of renal injury. In obesity specifically, the reduction of weight relieves obesity-related glomerular hyperfiltration and the associated glomerulopathy, adding a weight-mediated component on top of the direct effects.

14.3 Weight-independence and the place among renoprotective therapies

As with the cardiovascular benefit, the renal protection appears to be only partly attributable to weight loss — the haemodynamic, natriuretic and anti-inflammatory actions operate as more direct effects of GLP-1-receptor signalling, and albuminuria reductions have been observed relatively early. The practical importance mirrors the cardiovascular case: a patient whose weight response is modest may still gain renal protection, because the mechanism is not solely weight-dependent. This positions semaglutide as a distinct pillar alongside the established renoprotective therapies. Renin–angiotensin blockade (ACE inhibitors or ARBs), SGLT2 inhibitors, and the non-steroidal mineralocorticoid-receptor antagonist finerenone each protect the kidney through overlapping but non-identical mechanisms, and semaglutide adds a further, largely complementary mechanism. The emerging model of diabetic kidney disease management is therefore combination therapy across these pillars in high-risk patients, with semaglutide contributing both renal and cardiovascular protection.

14.4 Clinical implications

The renal data have direct practical consequences. For a patient with type 2 diabetes and chronic kidney disease, semaglutide now offers evidence-based slowing of renal progression and should be considered as part of the renoprotective regimen alongside RAS blockade, an SGLT2 inhibitor and, where appropriate, finerenone. The pharmacokinetics support this use: because semaglutide is cleared by metabolism rather than by renal excretion of intact drug, *no dose adjustment is required* across the range of renal function (Chapter 8), and the drug could be studied in, and benefits, the CKD population. Two practical cautions apply. First, gastrointestinal adverse effects that cause vomiting or diarrhoea can precipitate dehydration and a pre-renal fall in kidney function, so hydration and dose management during the titration phase matter especially in patients with existing renal impairment (Chapter 17). Second, as glucose falls, the doses of other glucose-lowering agents (particularly insulin and sulfonylureas) should be reviewed, as renal impairment itself alters their clearance. Overall, the renal indication extends semaglutide's role from metabolic and cardiovascular protection

15 OUTCOMES III

Liver, Heart Failure and Sleep Apnoea

The organ-specific benefits of this chapter are, in a sense, the most direct demonstration of the thesis of Chapter 6 — that these conditions are downstream expressions of the same ectopic-fat and inflammatory phenotype, and that correcting the phenotype treats them. Metabolic liver disease, obesity-related heart failure with preserved ejection fraction, and obstructive sleep apnoea are three organ diseases of adiposity, and semaglutide's effect on each follows from removing the fat and inflammation that drive them. This chapter reviews the evidence in each and the mechanism that connects them.

15.1 Liver — MASLD / MASH

Metabolic dysfunction-associated steatotic liver disease (MASLD) and its progressive form, steatohepatitis (MASH), are the hepatic face of ectopic fat: lipid accumulating in the hepatocyte (Chapter 5) drives inflammation (steatohepatitis) and, over time, fibrosis and cirrhosis. MASH has become a leading cause of chronic liver disease and liver transplantation, and until recently had almost no effective pharmacotherapy. The **ESSENCE** programme tested semaglutide in biopsy-confirmed MASH with fibrosis and demonstrated substantially higher rates of **resolution of steatohepatitis** — on the order of 63% versus about 34% with placebo — a large and clinically meaningful effect. The mechanism is exactly the ectopic-fat depletion of Chapter 5 applied to the liver: reducing hepatic lipid removes the substrate for lipotoxic injury and inflammation, and the drug's systemic anti-inflammatory and insulin-sensitising actions reinforce the effect. An honest nuance is that resolution of the *steatohepatitis* (the inflammatory activity) is more readily achieved than regression of established *fibrosis* (the scarring), which responds more slowly and less completely; nonetheless, halting the inflammatory driver is expected to slow or prevent fibrosis progression, and the emerging hepatic indication (Chapter 9) reflects a genuine advance in a disease that badly needed one.

15.2 Heart failure with preserved ejection fraction

Heart failure with preserved ejection fraction (HFpEF) is a heterogeneous syndrome, but a large and increasingly recognised subset is fundamentally *obesity-related* — driven by the volume expansion, systemic inflammation, epicardial fat, and haemodynamic burden that obesity imposes on the heart. The **STEP-HFpEF** trial tested semaglutide 2.4 mg specifically in this obesity-related HFpEF phenotype and showed clinically important improvements in **heart-failure symptoms and physical limitation** (measured by the Kansas City Cardiomyopathy Questionnaire), in **exercise capacity** (six-minute walk distance), and in weight and inflammatory markers (CRP). The mechanism is coherent with the rest of the monograph: by reducing weight, volume, visceral and epicardial fat, and systemic inflammation, semaglutide relieves the specific pathophysiology that obesity imposes on the failing-but-preserved ventricle. This reframes a subset of HFpEF as a treatable consequence of obesity rather than an intractable syndrome, and it positions semaglutide as disease-modifying for those patients — a striking example of an appetite-and-metabolism drug improving a cardiac syndrome by addressing its upstream cause.

15.3 Sleep apnoea and other systems

Obstructive sleep apnoea (OSA) is another mechanical disease of adiposity — fat deposition around the pharyngeal airway and the effects of obesity on respiratory mechanics narrow and destabilise the upper airway during sleep. Weight loss has long been known to improve OSA, and the weight loss produced by semaglutide improves it accordingly; however, the *dedicated drug indication* for OSA belongs to the dual agonist tirzepatide (SURMOUNT-OSA, developed in its companion volume), which was the first pharmacological agent approved specifically for the condition. For semaglutide, OSA improvement is a welcome consequence of weight loss rather than a formal indication. Beyond these three organ systems, the reach of the integrated phenotype extends further: weight loss and improved insulin sensitivity benefit conditions such as polycystic ovary syndrome (improving ovulation and metabolic features), osteoarthritis-related disability, and the metabolic complications of obesity generally. The unifying principle is the one established in Chapter 6 — that a therapy which corrects the underlying adiposity and inflammation improves the many organ-specific diseases that flow from them, which is why the list of benefits keeps lengthening as trials are done.

◆ KEY POINTS — CHAPTER 15

- **Liver (MASH):** ESSENCE showed steatohepatitis resolution ~63% vs ~34% — depleting hepatic ectopic fat (Chapter 5) removes the lipotoxic/inflammatory driver; resolution of inflammation precedes and outpaces regression of fibrosis.
- **HFpEF:** STEP-HFpEF improved heart-failure symptoms (KCCQ), exercise capacity and inflammation in obesity-related HFpEF — reframing a subset of HFpEF as a treatable consequence of obesity.
- **OSA:** weight loss improves sleep apnoea, but the dedicated OSA indication belongs to tirzepatide (SURMOUNT-OSA); for semaglutide it is a benefit of weight loss, not a formal indication.
- Other benefits (PCOS, osteoarthritis-related disability, and more) follow the same principle: correcting the adiposity/inflammation phenotype treats its many organ-specific expressions.

16 SAFETY — THE EYE

Diabetic Retinopathy and the NAION Signal

The eye deserves its own chapter because two distinct ocular issues are frequently conflated — one a treatable consequence of therapeutic *success*, the other a rare ischaemic event of uncertain causal link. Keeping them separate is essential both for counselling patients accurately and for taking the right preventive action, and this chapter develops the mechanism and the clinical approach for each.

16.1 Early worsening of diabetic retinopathy

SUSTAIN-6 reported more retinopathy complications with semaglutide (hazard ratio around 1.8), concentrated in patients with **pre-existing retinopathy, long-standing poorly controlled diabetes, and the largest, fastest HbA1c falls**. This is the well-described phenomenon of **early worsening**: the rapid normalisation of chronic hyperglycaemia transiently destabilises the retinal microvasculature before the long-term benefit of good control accrues. Crucially, it is a consequence of *rapid glycaemic improvement* itself, not a unique toxin of semaglutide — the same early worsening was described decades ago with intensive insulin therapy in the landmark diabetes-control trials, where tight control produced a transient worsening of retinopathy followed by substantial long-term benefit. Recognising early worsening as a general phenomenon of rapid glucose lowering, rather than a drug-specific effect, is the key to responding to it correctly: it is largely predictable and largely preventable.

16.2 The mechanism of early worsening

Understanding why rapid glucose lowering transiently worsens retinopathy clarifies who is at risk and how to prevent it. The retina of a chronically hyperglycaemic patient has adapted to its abnormal metabolic and haemodynamic environment; sudden correction perturbs several systems at once. Retinal blood flow, which is altered in chronic hyperglycaemia, shifts as glucose falls, and the autoregulation of the compromised retinal vasculature copes poorly with the change. Growth-factor signalling is also perturbed: rapid metabolic change can transiently alter levels of 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. The result is that eyes with *established* retinopathy — which already have the vascular fragility and neovascular tendency for these perturbations to act upon — can worsen in the first months of rapid improvement, before the durable benefit of lower glucose on the microvasculature takes hold. Eyes without significant baseline retinopathy have little for the phenomenon to act on and are at low risk. This mechanism explains precisely the risk profile seen in the trials — the danger is concentrated in the poorly controlled patient with existing retinopathy undergoing a large, fast fall in HbA1c — and it points directly to the preventive strategy.

PEARL In a patient with diabetes, poor control and known retinopathy, arrange a baseline dilated fundus exam, lower glucose more gradually, and co-manage with ophthalmology. The patient at risk is identifiable in advance.

16.3 NAION — a rare ischaemic optic neuropathy

Non-arteritic anterior ischaemic optic neuropathy (NAION) presents as sudden, usually painless, monocular vision loss from impaired perfusion of the optic-nerve head; the visual loss is generally irreversible. It is an entirely different entity from retinopathy — an ischaemic event of the optic nerve rather than a microvascular disease of the retina — and its possible association with semaglutide emerged more recently and remains genuinely uncertain. Following a 2024 single-centre signal, large studies gave mixed results: several cohorts and a meta-analysis found roughly a **two-fold** association (pooled hazard ratio around 2.6, becoming statistically significant mainly after about two years of exposure), while other matched analyses found **no** excess. 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), corresponding to roughly one extra case per 10,000 person-years, and recommended stopping semaglutide if NAION is confirmed; professional ophthalmology societies judged the absolute risk low and the data largely observational.

The interpretive difficulty is a textbook case of **confounding by indication**. The identified risk factors for NAION — older age, male sex, long diabetes duration, high HbA1c, existing diabetic retinopathy, obesity, and a small, "crowded" optic disc — overlap almost entirely with the population that takes semaglutide. When the people at highest baseline risk of an event are also the people most likely to receive a drug, an association will appear in observational data whether or not the drug is causal, and disentangling the two requires careful design that the current evidence has only partly achieved. Whether the signal reflects a true, mechanistically plausible effect (perhaps a haemodynamic or perfusion change at a vulnerable optic disc, possibly related to rapid metabolic shift) or is largely an artefact of confounding is not yet resolved. What is clear is that the *absolute* risk is very low, that the association is at most modest, and that the appropriate response is proportionate awareness rather than alarm.

Table 16.1. Two ocular issues, contrasted.

	Early retinopathy worsening	NAION
Driver	Rapid HbA1c fall	Optic-nerve-head ischaemia (uncertain causal link)
Reversible?	Often transient	Usually permanent
Frequency	Uncommon, predictable	Very rare
Prevention/ action	Baseline exam, gradual control, ophthalmology	Counsel; urgent review and stop drug if sudden vision loss

17 SAFETY IN DEPTH

Adverse Effects — Mechanisms, Predictors, Prevention and Confounders

Adverse effects become manageable once you see that most trace to **one of two mechanisms**: the same delayed gastric emptying and area-postrema signalling that drives benefit, or the consequences of **rapid weight loss itself**. A third, quieter theme runs through the chapter — 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 activation of the area postrema (the brain's emetic sensor, deliberately accessible to the drug).[33] They are intrinsic to the pharmacology, not idiosyncratic.

Predictors. Faster titration, higher dose, female sex, and the period immediately after each dose increase. Tirzepatide carries the highest vomiting signal of the class; semaglutide and orforglipron are similar to each other.[33]

Prevention. Slow titration is the single best tool; eat small, low-fat, low-volume meals; stop at fullness; maintain hydration. If symptoms are marked, **hold or step down** a dose rather than pushing on, and allow extra weeks at the tolerated dose. Most GI effects fade with time at a stable dose.

Time course and escalation of management. The predictable trajectory is that nausea peaks in the days after each dose increase and settles over the subsequent weeks as gastric-emptying tachyphylaxis develops (Chapter 6), so the symptom pattern is one of transient peaks superimposed on a declining baseline rather than a steady burden. This trajectory is itself reassuring to convey to patients, because it reframes early nausea as a self-limiting adaptation rather than an indefinite penalty. 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. Vomiting and diarrhoea additionally matter because of their downstream consequences — dehydration and its renal effects (Section 17.7) and, for tirzepatide, reduced oral-contraceptive absorption (Chapter 8) — so their management is not merely about comfort. A minority of patients cannot tolerate the drug at any useful dose despite optimal measures, and recognising this early, rather than repeatedly forcing escalation, spares them a miserable and futile course.

17.2 Gallbladder and biliary disease

Why. Two mechanisms combine: **rapid weight loss** mobilises cholesterol and supersaturates bile (a risk with weight loss of any cause, including surgery), and reduced gallbladder **contractility** promotes bile stasis.[35] **Predictors:** rapid or large weight loss, higher dose, longer use, female sex, prior gallstones. **Prevention:** avoid excessively rapid loss; counsel on biliary symptoms; investigate right-upper-quadrant pain promptly.

17.3 Pancreatitis — signal, and the confounding problem

A 2025 meta-analysis of 62 randomised trials (≈ 66,000 patients) found a modestly increased relative risk of acute pancreatitis (≈ 1.4) with **no** association with pancreatic cancer.[34] Crucially, the patients who qualify for these drugs **already carry pancreatitis risk factors** — gallstones, hypertriglyceridaemia, obesity, alcohol use, diabetes — and rapid weight loss is itself a risk factor. So part of the observed signal is the population, not purely the molecule. **Action:** avoid in those with a clear prior history where alternatives exist; stop and investigate if severe persistent abdominal pain occurs; do not restart after confirmed drug-related pancreatitis.

17.4 Pulmonary aspiration and anaesthesia

Why. Delayed gastric emptying means **retained gastric contents despite standard fasting**; aspiration of undigested food has occurred at induction of anaesthesia and during endoscopy.[31,32] **Predictors:** active dose escalation, current GI symptoms, higher dose. **Prevention:** follow anaesthesia-society guidance to hold the agonist before elective procedures (typically holding the weekly dose the week of surgery), consider point-of-care gastric ultrasound, prolong clear-fluid fasting, and treat as "full stomach" if in doubt. Note the confounder: residual gastric contents are also more common in long-standing diabetic gastroparesis *without* any GLP-1 agonist.

17.5 Loss of lean (muscle) mass

Any large weight loss includes lean as well as fat mass; the concern is functional sarcopenia, especially in older patients. **Prevention** is behavioural, not pharmacological: protein ≈ 1.2–1.6 g/kg/day and **resistance training** throughout treatment (Chapter 18). Agents that preserve lean mass are in development.[43]

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-level data have **not** demonstrated an increase in medullary thyroid carcinoma. The contraindication remains precautionary in MTC/MEN-2 families.

Suicidality / psychiatric. Early pharmacovigilance flags prompted large analyses that 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, not a direct nephrotoxin) — maintain hydration, especially with intercurrent illness. **Euglycaemic ketoacidosis** has been reported, particularly when co-prescribed with an SGLT2 inhibitor or during very low carbohydrate intake. Injection-site reactions are mild. For the eye, see Chapter 16.

18 PRACTICE

Practical Management — Diet, Muscle, Plateau and Discontinuation

Semaglutide is best understood as a tool that makes durable lifestyle change *possible*, not as a substitute for it. By quieting hunger and food noise (Chapters 6 and 7), the drug removes the biological headwind that defeats most dietary efforts — but the *quality* of the weight lost, the preservation of muscle and function, the handling of plateaus, and what happens on stopping all depend on how the treatment period is managed. This chapter is the practical counterpart to the mechanism chapters: it translates the pharmacology into the diet, training, monitoring and long-term strategy that determine whether a patient ends up healthier or merely lighter. The central message is that the drug and an active nutritional-and-training programme are complementary, and that the reduced-intake window the drug creates is an opportunity to build better habits, not a reason to neglect them.

18.1 Diet and protein

Because appetite is suppressed, patients on semaglutide eat considerably less, and this makes the *composition* of what they do eat far more important than before — a smaller intake must still meet the body's needs. Two priorities dominate. The first is **adequate protein**. During weight loss, sufficient dietary protein is the primary nutritional defence of lean (muscle) mass, and a reasonable target is on the order of **1.2–1.6 g per kg of body weight per day** (using a sensible reference or goal weight in those with severe obesity), distributed across meals rather than concentrated in one, and favouring high-quality protein sources that supply the full range of amino acids. Because total intake is small, hitting the protein target often requires deliberate planning — protein is prioritised at each meal, and protein supplementation is sometimes useful when appetite is most suppressed. The second priority is **nutrient density**: with fewer total calories, the diet must be rich in vegetables, fruit, whole grains, and sources of micronutrients, so that the reduced intake does not create nutritional gaps. Attention to adequate fibre helps the constipation that can accompany reduced intake, and attention to hydration is important both for gastrointestinal comfort and to protect renal function during any vomiting or diarrhoea (Chapter 17). Practically, the same measures that support nutrition also ease gastrointestinal tolerability — smaller, lower-fat, unhurried meals — so good dietary practice and good tolerability reinforce each other. Where intake is very restricted or prolonged, attention to specific micronutrients (including vitamin B12 and others) is reasonable, particularly in patients with additional risk factors such as prior bariatric surgery.

18.2 Resistance training preserves muscle

This section carries one of the most important practical messages of the monograph, because it addresses the chief downside of large weight loss. Any substantial weight loss — by diet, drug or surgery — includes a component of **lean (muscle) mass** alongside fat, and the loss of muscle carries functional and metabolic costs: reduced strength and mobility, a lower resting metabolic rate (which itself predisposes to later regain), and, in older patients especially, the risk of clinically significant **sarcopenia** and frailty. The goal of good management is therefore

not merely weight loss but favourable *body recomposition* — losing fat while preserving, or even building, muscle.

The intervention that preserves lean mass is **resistance (strength) training**, and it is not optional for anyone who wishes to protect their muscle and function through a large weight loss. A practical programme involves progressive resistance exercise on the order of two to three sessions per week, working the major muscle groups through a full range of motion with a load and volume that are gradually increased over time (the principle of progressive overload), begun early in treatment and continued throughout. Resistance training works synergistically with the protein intake of Section 18.1 — protein provides the substrate and the training provides the stimulus for muscle maintenance and growth — so the two must be pursued together rather than in isolation. Aerobic (cardiovascular) exercise has its own substantial benefits for cardiovascular and metabolic health and should also be encouraged, but it is *resistance* training specifically that defends muscle mass, and it should not be displaced by cardio alone. For older patients, in whom the sarcopenia risk is greatest and the functional stakes highest, the priority on resistance training and adequate protein is correspondingly greater, not lesser. Pharmacological agents designed specifically to preserve or build lean mass during weight loss are in development (Chapter 19) and may eventually complement this approach, but for now the behavioural combination of resistance training plus adequate protein is the established and effective strategy.

18.3 The plateau

A weight-loss plateau is an expected physiological equilibrium, not a failure, and it is analysed in detail in Chapter 11. In brief, as weight falls, energy expenditure declines and the defended set-point partially reasserts itself, so loss eventually ceases at a new equilibrium. The practical response depends on whether the patient is at the maximum dose: if not, and the current dose is tolerated, escalation to the next step is the primary lever (having first excluded a pseudo-plateau from incomplete titration or a dose not yet at steady state); if already at the maximum dose, the levers are behavioural and strategic — reinforcing the diet and resistance-training programme of this chapter, verifying adherence and administration technique, reconsidering whether the plateau is an acceptable maintenance weight, and, where greater effect is genuinely needed, considering a switch to a higher-efficacy agent. The practical management chapter and the dosing chapter thus meet at the plateau: the dosing lever and the lifestyle lever are used together, and neither substitutes for the other.

18.4 Discontinuation and weight regain

The most important long-term counselling point, and one that must be discussed *before* starting rather than at the moment of stopping, is that discontinuation predictably leads to **weight regain**. The withdrawal trials (STEP-4, and SURMOUNT-4 for the dual agonist) showed that stopping the drug after successful loss led to regain of a large proportion of the lost weight over the following year, while continuation maintained the loss (Chapter 12). This is not a failure of will or of the drug but the direct consequence of the set-point physiology of Chapter 6: the drug *overrides* the defended body-weight set-point while it is present, and removing it allows the set-point to reassert itself, restoring appetite and, with it, the weight and the ectopic fat and the metabolic risk. The clinical implications are substantial. First, obesity should be framed and treated as a **chronic, relapsing disease** requiring long-term

19 LANDSCAPE

The GLP-1 Landscape and Future Directions (2026)

Semaglutide sits in a field moving quickly toward greater efficacy, oral delivery and broader indications. This closing chapter places it in the landscape of 2026, surveys the agents that are newer or coming, and identifies the directions of travel and the questions that remain — before drawing together the thesis that has run through the whole monograph.

19.1 Where semaglutide stands

Semaglutide remains the most *evidence-rich single agent* in the field, with mature outcome data spanning glycaemia, weight, cardiovascular disease (SELECT), kidney disease (FLOW) and, increasingly, liver disease (ESSENCE), and it is now available in both subcutaneous and oral forms. What it is *not*, any longer, is the most potent agent for weight loss: the dual agonist tirzepatide exceeds it (Chapter 12), and the investigational triple agonist retatrutide exceeds tirzepatide in early trials. Semaglutide's enduring strengths are therefore the breadth and maturity of its hard-outcome evidence, its oral availability, and its established safety record — the qualities that, for many patients and indications, matter more than the last few percentage points of weight loss.

19.2 The 2026 landscape

Table 19.1 summarises the agents and classes that define the field as of 2026, from the newly available oral options to the higher-efficacy multi-agonists in development.

Table 19.1. The 2026 incretin landscape.

Agent / class	Status (2026)	Weight effect (approx.)
Oral semaglutide (weight management)	Approved late 2025	Somewhat below injectable
Oral semaglutide tablets (T2D)	Higher-dose tablets, early 2026	—
Tirzepatide (dual GIP/GLP-1)	Established; exceeds semaglutide on weight	≈ 20–21%
Orforglipron (oral small-molecule GLP-1)	Approved Apr 2026 (weight); needle-free, scalable	≈ 12%
Retatrutide (triple GIP/GLP-1/ glucagon)	Phase 3 (TRIUMPH); projected ~2028	≈ 24% (phase 2)
CagriSema / amylin combinations	Late-stage development	High

19.3 Oral and small-molecule delivery

One of the two defining trends is the move away from injection. Oral semaglutide (Chapter 2) proved that a peptide could be given by mouth at all, and the arrival of oral formulations for

weight management in late 2025 and higher-dose oral tablets for diabetes in early 2026 extended that option. The more transformative development, however, is the emergence of **orforglipron**, the first oral *non-peptide, small-molecule* GLP-1 receptor agonist, approved in April 2026 for chronic weight management. Because it is a small molecule rather than a peptide, orforglipron is far easier and cheaper to manufacture at scale and does not require the SNAC-dependent gastric absorption window or the cold chain of the peptides — it is a conventional pill. Its weight-loss magnitude (around 12%) is below injectable semaglutide, but its scalability and convenience could dramatically widen access, particularly in settings where cost and cold-chain logistics have limited the peptides. The small-molecule route thus points toward a future in which effective incretin therapy is as accessible as an ordinary daily tablet.

19.4 Higher-efficacy multi-agonism

The second defining trend is the pursuit of greater efficacy through the recruitment of additional hormonal pathways — the strategy that produced tirzepatide and is being extended further. The dual GIP/GLP-1 agonist tirzepatide is established (its companion volume develops it in full). Beyond it, the triple agonist **retatrutide** adds a *glucagon-receptor* arm to GIP and GLP-1, and produced weight loss of around 24% in phase 2, with phase 3 (the TRIUMPH programme) underway and projected availability toward 2028. The rationale for the glucagon arm illustrates the sophistication of the approach: 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 agonists exploit those benefits while relying on their potent GLP-1 (and GIP) activity to keep glycaemia controlled (a point developed in Chapter 4). The trajectory is clear: progressively greater efficacy by layering complementary hormonal actions on the GLP-1 foundation, moving medical therapy ever closer to — and, for some agents, into — the weight-loss range previously reached only by surgery.

19.5 Other mechanisms — amylin and muscle preservation

Two further directions deserve mention. The first is the **amylin** axis: amylin analogues (such as cagrilintide, including in combination with semaglutide as "CagriSema") act through a different satiety pathway and, in combination with incretin agonism, produce high degrees of weight loss, representing a distinct hormonal lever being added to the field. The second, and particularly relevant to the practical concerns of Chapter 18, is the development of **muscle-preserving agents** — therapies (for example, those targeting the myostatin/activin pathway) intended to preserve or increase lean mass during weight loss, so that the loss is more purely fat. If such agents prove effective and safe, they could be combined with incretin therapy to address the lean-mass loss that is currently countered only behaviourally (through protein and resistance training), improving the *quality* of weight loss rather than merely its magnitude. Together, the amylin and muscle-preservation directions show that the field is diversifying beyond incretins alone toward a broader pharmacology of energy balance and body composition.