The Policymaker's Bias: How Cognitive Failures Block Rational Responses to GLP-1 Obesity Medications

Pablo M. Zavala-Reina

Heinz College of Public Policy and Management

Carnegie Mellon University

Author Note

Working manuscript. Submitted March 24, 2026. Last edited March 24, 2026.

Abstract

GLP-1 receptor agonists have changed the clinical frontier of obesity treatment faster than public policy has changed its governing assumptions. Semaglutide produces large mean weight reductions, tirzepatide pushes those reductions further, and semaglutide also lowers major adverse cardiovascular events in adults with overweight or obesity and cardiovascular disease. Yet Medicare still does not broadly cover anti-obesity drugs for obesity alone, Medicaid coverage remains fragmented, and prescribing remains disproportionately low relative to eligibility. This paper argues that the adoption gap is not best understood as a simple shortage of information. It is a behavioral-economics problem inside institutions. Policymakers are anchoring on a narrow fiscal score, inheriting a status-quo default designed for a pre-GLP-1 era, and responding to uncertainty with loss-averse caution. A compact formalization shows that the policy problem is multi-attribute, not one-dimensional: once mortality, long-run cost offset, equity, and adherence are brought back into the objective, the apparent rationality of exclusion weakens substantially. In that setting, a fiscally correct projection can still reverse the welfare ranking of policies, and ambiguity-averse institutions can rationally delay action even as the mean clinical case improves. The proposed intervention is twofold: default coverage with active opt-out at the institutional level, and loss-framed adherence communication at the patient level. The broader implication is that the state now faces a mirror-image version of the patient's own cognitive problem. In both settings, the central question is not whether evidence exists, but whether the choice architecture can preserve human priorities while using it well.

Keywords: behavioral economics, obesity policy, GLP-1 receptor agonists, public policy, choice architecture

The Policymaker's Bias: How Cognitive Failures Block Rational Responses to GLP-1 Obesity Medications

I. The Issue

GLP-1 receptor agonists have shifted the obesity debate from speculative promise to measurable clinical effect. In the STEP 1 trial, adults with overweight or obesity receiving semaglutide 2.4 mg experienced a mean body-weight reduction of 14.9% over 68 weeks (Wilding et al., 2021). In SURMOUNT-1, tirzepatide reached mean reductions as high as 22.5% (Jastreboff et al., 2022). The stronger policy fact is cardiovascular: in SELECT, semaglutide reduced the composite risk of cardiovascular death, nonfatal myocardial infarction, and nonfatal stroke by 20% in adults with preexisting cardiovascular disease and overweight or obesity but without diabetes (Lincoff et al., 2023). One U.S. cost-effectiveness model placed semaglutide near $23,000–$27,000 per quality-adjusted life year under its assumptions (N. Kim et al., 2022).

The policy response has not kept pace. Under current law, Medicare Part D can generally cover anti-obesity drugs only when they are used for a medically accepted indication other than obesity itself, such as diabetes or cardiovascular risk reduction (Cubanski & Williams, 2024). Medicaid policy remains fragmented: KFF reported that only 13 state fee-for-service Medicaid programs covered GLP-1s for obesity treatment as of January 2026 (Williams, 2026). Prescribing gaps are also large relative to clinical eligibility. In a national electronic-health-record study, fewer than 3% of eligible patients had received semaglutide or tirzepatide for obesity, and Hispanic and non-Hispanic Asian patients were less likely than White patients to receive those therapies (C. Kim et al., 2025). These access gaps sit on top of an unequal disease burden. Non-Hispanic Black adults had the highest obesity prevalence in the 2017–2018 CDC estimates (Hales, Carroll, Fryar, & Ogden, 2020).

This is not a marginal policy domain. The Lancet Commission on the global syndemic described excess body weight as a condition affecting billions of people and imposing macroeconomic costs measured in shares of gross domestic product (Swinburn et al., 2019).

Yet public institutions still treat obesity pharmacotherapy as a narrow pharmaceutical line item inside a policy world designed for episodic treatment and short budget windows.

This paper argues that the gap between pharmacological efficacy and policy adoption is not primarily a shortage of evidence. It is a failure of institutional choice architecture. The same behavioral tendencies that distort individual decision-making also distort legislative and administrative reasoning. The patient's brain and the policy system are not identical, but they rhyme. Both are vulnerable to salience, inertia, and distorted responses to uncertainty. Once that parallel is made explicit, the central policy question changes. The problem is no longer merely whether GLP-1 medications work. It is whether the state is evaluating them with the right objective function.

II. Behavioral Diagnosis

The first distortion is anchoring on a misspecified objective. The Congressional Budget Office estimates that authorizing Medicare coverage of anti-obesity medications would increase federal spending over the 2026–2034 window by about $35 billion (Congressional Budget Office, 2024). That estimate matters politically because it arrives as the first salient number in the debate. In the language of Tversky and Kahneman (1974), it functions as an anchor. But the deeper issue is structural. A proper policy evaluation is multi-attribute:

$$V(a_j) = \sum_{i=1}^{n} w_i v_i(x_{ij}), \tag{1}$$

where policies are evaluated not only on short-run fiscal cost, but also on mortality reduction, equity, persistence, and downstream savings. The CBO score is not useless. It is a projection of that broader problem onto one dimension. Within that restricted space, the estimate can be correct. Relative to the full policy problem, however, it can still be normatively incomplete. That distinction matters because obesity imposes large direct medical costs on the United States, including substantial excess spending per affected adult (Cawley et al., 2021). A government that scores only ten-year expenditure is answering a narrower question than the one it claims to care about.

The sharper claim is that projection can change the ranking itself. If F denotes the fiscal component and U the omitted attributes, then for any two policies a and b the relevant comparison is

$$\Delta_V(a, b) = \Delta_F(a, b) + \Delta_U(a, b). \tag{2}$$

If $\Delta_F(a, b) < 0$ but $\Delta_U(a, b) > -\Delta_F(a, b)$, the projected fiscal rule prefers a while the full welfare rule prefers b. Therefore, a fiscally disciplined score can still be a welfare-misspecified choice rule. That is the more rigorous statement of the anchoring problem.

The second distortion is status quo bias embedded in legislative defaults. Medicare exclusion is not a neutral baseline; it is an inherited default from an earlier regulatory era. Behavioral evidence on defaults is clear: automatic enrollment dramatically changes take-up without changing the underlying instrument (Madrian & Shea, 2001). The mechanism is inertia. Public policy is full of the same inertia: states must actively choose to cover GLP-1s for obesity, and Congress must actively reverse statutory exclusion. Therefore, noncoverage should be treated as an active policy setting, not as a natural baseline.

The third distortion is loss aversion under genuine uncertainty. The modern evidence base is unusually rich but not psychologically simplifying. The drugs work better than expected. The eligible population is large. Discontinuation erodes gains. Budget exposure is highly salient. New federal models such as BALANCE indicate that CMS is already experimenting with negotiated pricing and with pairing medication access to lifestyle support, including Medicaid launch as early as May 2026 and a $50 monthly payment for eligible Medicare beneficiaries in the bridge demonstration (Centers for Medicare & Medicaid Services, 2026). This is informative, but it does not collapse uncertainty. It redistributes it. Under prospect theory, losses loom larger than gains (Kahneman & Tversky, 1979). When policymakers face a choice with visible fiscal downside and delayed or distributed social upside, caution is not random. It is behaviorally predictable.

The relevant decision problem is not simple expected value maximization under a

single trusted distribution. A more realistic institutional criterion is therefore

$$W(a;\mathcal{P}) = \min_{p\in\mathcal{P}} \mathbb{E}_p[V(a,\omega)], \tag{3}$$

where ω indexes negotiated price, uptake, persistence, adverse events, and downstream fiscal offset, and $\mathcal{P}$ is a set of plausible posteriors rather than one posterior the institution fully trusts. Under that criterion, richer evidence can still slow adoption if it expands the set of plausible downside states faster than it sharpens the center of the distribution. Therefore, more evidence does not mechanically imply more willingness to cover.

A final distortion is moralization. Obesity still attracts a policy discourse saturated with ideas of weak will, poor character, and self-inflicted risk. That discourse is increasingly inconsistent with the evidence. Reviews of obesity neuroscience emphasize disruptions in reward and control circuitry rather than simple preference failure (Volkow, Wang, & Baler, 2011). Experimental work on GLP-1 signaling shows that these therapies modulate appetite- and reward-related brain areas in humans (van Bloemendaal et al., 2014). At the same time, international expert consensus treats weight stigma as a meaningful barrier to health and care (Rubino et al., 2020). The point is not that biology eliminates agency. It is that policy continues to moralize a condition whose mechanisms are increasingly being described in neurobehavioral rather than purely voluntarist terms. When a disease is moralized, policymakers not only underweight treatment. They also overestimate how much effort a patient should be expected to supply unaided, which makes noncoverage appear more defensible than it is.

III. The Intervention

The intervention has two linked parts, one institutional and one patient-facing. Institutionally, the paper proposes default coverage with active opt-out for FDA-approved anti-obesity medications that satisfy a cost-effectiveness screen and are paired with behavioral-support infrastructure. The behavioral logic is straightforward. If exclusion remains the default, fiscal cost enters as the most immediate loss. If coverage becomes the

default, foregone health enters as the more salient loss. In a prospect-theory frame, the physical policy can remain the same while the reference point changes:

$$u(x) = \begin{cases} x^{\alpha}, & x \geq 0, \\ -\lambda(-x)^{\alpha}, & x < 0, \end{cases} \tag{4}$$

with $\lambda > 1$. Therefore, the institutional intervention does not depend on changing every belief inside the system. It depends on changing which dimension is coded as the default loss.

The second intervention is patient-facing and addresses present bias. Gastrointestinal side effects are common in GLP-1 treatment, and treatment discontinuation is followed by substantial weight regain and reversal of cardiometabolic improvement (Wilding et al., 2021, 2022). Standard adherence communication often frames continuation as a gain. A behavioral alternative is to frame discontinuation as a loss: stopping treatment puts already-earned cardiovascular and metabolic gains at risk. This is a more behaviorally realistic translation of the temporal trade-off patients face.

The bundle argument also has a dynamic microfoundation. Let s_t denote persistence, d_t medication access, b_t behavioral support, and c_t side-effect burden. A compact state equation is

$$s_{t+1} = \rho s_t + \eta b_t - \gamma c_t, \qquad h_{t+1} = h_t + \beta d_t s_t, \tag{5}$$

with $0 \leq \rho < 1$ and $\eta, \gamma, \beta > 0$. In that representation, behavioral support does not merely add encouragement in parallel to medication. It acts on the persistence channel through which medication generates cumulative health gains. Hence, the policy value of support is partly mediated, not only additive.

Bundling matters here. The STEP 3 trial showed that semaglutide plus intensive behavioral therapy produced larger average weight loss than behavioral therapy alone (Wadden et al., 2021). That arithmetic does not, by itself, prove supermodularity for every outcome. But it is enough to motivate a broader hypothesis: drug access and behavioral support act through different channels and should not be evaluated as independent policy

levers. In the policy setting, a reasonable formal condition is

$$\frac{\partial^2 H}{\partial d\,\partial b} \geq 0, \tag{6}$$

where H is a long-run health outcome, d is drug access, and b is behavioral support. If that cross-partial is positive for survival, persistence, or cardiometabolic durability, then piecemeal provision systematically understates the value of the bundle.

IV. Implementation and Evaluation

The institutional intervention should first be tested in a randomized survey experiment with state Medicaid directors and senior health-policy officials. Stratification by political context and current coverage status would help isolate whether framing effects travel across partisan settings or only within friendly policy environments. The treatment arm would receive a default-approval frame: coverage is provisionally authorized unless the state opts out with a public justification discussing expected health consequences. The control arm would receive a standard opt-in frame in which the state must affirmatively add coverage while viewing the same fiscal information. The primary outcome would be stated willingness to cover. A secondary outcome would ask participants to allocate weights across policy attributes. That second measure matters because the theory is not merely that people say different things under different frames. It is that they attend to different dimensions of the problem.

The patient-level intervention is naturally evaluated through a parallel-arm randomized trial among new GLP-1 initiators in safety-net clinical settings. The treatment arm receives loss-framed adherence messages at predetermined intervals; the control arm receives conventional gain-framed encouragement. The primary endpoint is medication persistence at six months, with secondary outcomes for risk understanding and perceived side-effect tolerance. Stratification by race, income, and baseline body mass index is important because the policy goal is not only higher persistence in the abstract. It is a persistence intervention that remains credible across the populations most exposed to current

access disparities.

Formally, the causal estimands are simple even if the policy problem is not:

$$\tau_{\text{inst}} = \mathbb{E}[Y_i(1) - Y_i(0)], \qquad \tau_{\text{pat}} = \mathbb{E}[P_i(1) - P_i(0)], \tag{7}$$

where Y_i is the institutional coverage response and P_i is patient persistence. Random assignment identifies the framing effects. The harder question is therefore external validity: whether those effects survive translation from the experimental environment into durable administrative and clinical routines.

The two evaluation designs also speak to each other. If the official-level experiment changes which policy attributes are foregrounded, and the patient-level experiment changes how near-term discomfort is weighed against long-term health, then the package offers a behavioral intervention at two levels of analysis. The individual patient and the public official are both making intertemporal decisions under uncertainty. That is why the paper treats the institutional and patient arms as complements rather than isolated experiments.

V. Costs, Benefits, Limitations, and Forward Link

The case for reform should not be exaggerated. The fiscal problem is real. Even if one accepts the strongest clinical evidence, a large expansion of public coverage creates immediate budget exposure, and the political system discounts benefits that arrive later or outside the budget window. That is precisely why the current debate is so vulnerable to anchoring. Yet the opposite exaggeration is equally misleading. A narrow budget score is not the same thing as a complete social evaluation. A model can still show semaglutide as cost-effective under specified assumptions (N. Kim et al., 2022) while public payers remain worried about short-run affordability. Both statements can be true at once.
Cost-effectiveness asks whether a treatment generates enough value per unit of health gain. Affordability asks whether a payer can absorb the budget shock on the relevant timeline. Modern obesity policy repeatedly collapses those two questions into one, and that collapse is one more form of misspecification.

Three limitations deserve emphasis. First, stated-preference experiments are not enacted policy. Second, drug pricing remains a structural constraint that no amount of behavioral reframing can erase. Third, the bundle argument is strongest when states also invest in support architecture that makes long-run adherence plausible. A default can reorient choice, but it cannot substitute for food policy, urban design, or income support.

Even with those limits, the core diagnosis remains. The policy response to GLP-1 obesity medications is being filtered through an institutional mind that overweights the most immediate number, inherits yesterday's default, and reacts to uncertainty by protecting itself against visible losses. That is a failure of architecture. Therefore, the next step is not only better evidence, but better systems for organizing evidence across brains, agencies, and communities of interpretation. What is needed is an infrastructure that can keep long-horizon health effects, fiscal exposure, and distributional consequences in view at the same time, and make those trade-offs legible to both human decision-makers and the analytic tools that advise them. More generally, this is the class of problem that motivates parallel work on expert-editable, model-agnostic analytic systems in which objectives, evidence graphs, and policy menus remain explicit state objects rather than being flattened into prose after the fact. The objective must remain human, but the evidence space is now too large to manage well with unaided institutional attention. As public life enters an era in which human institutions will increasingly deliberate alongside AI agents and machine-mediated systems, the normative task is to ensure that those systems keep the right quantities in view: health, time, dignity, and fair distribution of opportunity.

References

Cawley, J., Biener, A., Meyerhoefer, C., Ding, Y., Zvenyach, T., Smolarz, B. G., & Ramasamy, A. (2021). Direct medical costs of obesity in the united states and the most populous states. *Journal of Managed Care & Specialty Pharmacy*, *27*(3), 354–366. doi: 10.18553/jmcp.2021.20410

Centers for Medicare & Medicaid Services. (2026). *Balance model: Better approaches to lifestyle and nutrition for comprehensive health.* Centers for Medicare & Medicaid Services. Retrieved from https://www.cms.gov/priorities/innovation/innovation-models/balance (Model overview and FAQ page accessed March 24, 2026)

Congressional Budget Office. (2024, October). *How would authorizing medicare to cover anti-obesity medications affect the federal budget?* (Tech. Rep.). Congressional Budget Office. Retrieved from https://www.cbo.gov/publication/60816

Cubanski, J., & Williams, E. (2024, November 26). *Proposed coverage of anti-obesity drugs in medicare and medicaid would expand access to millions of people with obesity.* KFF. Retrieved from https://www.kff.org/medicare/issue-brief/proposed-coverage-of-anti-obesity-drugs-in-medicare-and-medicaid-would-expand-access-to-millions-of-people-with-obesity/ (Policy watch article)

Hales, C. M., Carroll, M. D., Fryar, C. D., & Ogden, C. L. (2020). *Prevalence of obesity and severe obesity among adults: United states, 2017–2018* (Tech. Rep. No. 360). National Center for Health Statistics. Retrieved from https://www.cdc.gov/nchs/products/databriefs/db360.htm

Jastreboff, A. M., Aronne, L. J., Ahmad, N. N., Wharton, S., Connery, L., Alves, B., ... Stefanski, A. (2022). Tirzepatide once weekly for the treatment of obesity. *New England Journal of Medicine*, *387*(3), 205–216. doi: 10.1056/NEJMoa2206038

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*(2), 263–291. doi: 10.2307/1914185

Kim, C., Ross, J. S., Jastreboff, A. M., Meeker, D., Freedman, H., Krumholz, H. M., & Lu, Y. (2025). Uptake of and disparities in semaglutide and tirzepatide prescribing for obesity in the us. *JAMA*, *333*(24), 2203–2206. doi: 10.1001/jama.2025.4735

Kim, N., Wang, J., Burudpakdee, C., Song, Y., Ramasamy, A., Xie, Y., ... Sullivan, S. D. (2022). Cost-effectiveness analysis of semaglutide 2.4 mg for the treatment of adult patients with overweight and obesity in the united states. *Journal of Managed Care & Specialty Pharmacy*, *28*(7), 740–752. doi: 10.18553/jmcp.2022.28.7.740

Lincoff, A. M., Brown-Frandsen, K., Colhoun, H. M., Deanfield, J., Emerson, S. S., Esbjerg, S., ... Ryan, D. H. (2023). Semaglutide and cardiovascular outcomes in obesity without diabetes. *New England Journal of Medicine*, *389*(24), 2221–2232. doi: 10.1056/NEJMoa2307563

Madrian, B. C., & Shea, D. F. (2001). The power of suggestion: Inertia in 401(k) participation and savings behavior. *The Quarterly Journal of Economics*, *116*(4), 1149–1187. doi: 10.1162/003355301753265543

Rubino, F., Puhl, R. M., Cummings, D. E., Eckel, R. H., Ryan, D. H., Mechanick, J. I., ... Dixon, J. B. (2020). Joint international consensus statement for ending stigma of obesity. *Nature Medicine*, *26*(4), 485–497. doi: 10.1038/s41591-020-0803-x

Swinburn, B. A., Kraak, V. I., Allender, S., Atkins, V. J., Baker, P. I., Bogard, J. R., ... Dietz, W. H. (2019). The global syndemic of obesity, undernutrition, and climate change: The lancet commission report. *The Lancet*, *393*(10173), 791–846. doi: 10.1016/S0140-6736(18)32822-8

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, *185*(4157), 1124–1131. doi: 10.1126/science.185.4157.1124

van Bloemendaal, L., IJzerman, R. G., ten Kulve, J. S., Barkhof, F., Konrad, R. J., Drent, M. L., ... Diamant, M. (2014). Glp-1 receptor activation modulates appetite- and reward-related brain areas in humans. *Diabetes*, *63*(12), 4186–4196. doi: 10.2337/db14-0849

Volkow, N. D., Wang, G.-J., & Baler, R. D. (2011). Reward, dopamine and the control of food intake: Implications for obesity. *Trends in Cognitive Sciences*, *15*(1), 37–46. doi: 10.1016/j.tics.2010.11.001

Wadden, T. A., Bailey, T. S., Billings, L. K., Davies, M., Frias, J. P., Koroleva, A., ... Garvey, W. T. (2021). Effect of subcutaneous semaglutide vs placebo as an adjunct to intensive behavioral therapy on body weight in adults with overweight or obesity: The step 3 randomized clinical trial. *JAMA*, *325*(14), 1403–1413. doi: 10.1001/jama.2021.1831

Wilding, J. P. H., Batterham, R. L., Calanna, S., Davies, M., Van Gaal, L. F., Lingvay, I., ... Kushner, R. F. (2021). Once-weekly semaglutide in adults with overweight or obesity. *New England Journal of Medicine*, *384*(11), 989–1002. doi: 10.1056/NEJMoa2032183

Wilding, J. P. H., Batterham, R. L., Davies, M., Van Gaal, L. F., Kandler, K., Konakli, K., ... Kushner, R. F. (2022). Weight regain and cardiometabolic effects after withdrawal of semaglutide: The step 1 trial extension. *Diabetes, Obesity and Metabolism*, *24*(8), 1553–1564. doi: 10.1111/dom.14725

Williams, E. (2026, January 16). *Medicaid coverage of and spending on glp-1s.* KFF. Retrieved from https://www.kff.org/medicaid/medicaid-coverage-of-and-spending-on-glp-1s/ (Issue brief)