

The Complexity Laffer Curve: Optimal Menu Size within a Parametric Menu Family

Siddharth Karuturi

Illinois Mathematics and Science Academy

skaruturi@imsa.edu

Abstract

We study a monopolist's optimal choice of menu size when consumers have heterogeneous cognitive abilities. Within a parametric family of quality-price screening menus indexed by the number of tiers N , the monopolist commits to complexity anticipating that cognitively constrained consumers may exit, evaluate only a subset of options, or choose suboptimally. The interaction is formalized as a Stackelberg game. Our central result—the **Complexity Laffer Curve**—establishes that expected profit attains an interior maximum in N : screening gains from finer menus eventually give way to behavioral losses as cognitive overload erodes participation and decision quality. We prove existence of an interior profit-maximizing menu size N^* under explicit sufficient conditions, without claiming the profit profile is globally unimodal. Under a tractable exit-only variant, we establish three sharper results: constrained consumers yield weakly lower profit than the rational benchmark (Condition R); a first-order stochastic dominance (FOSD) shift toward higher cognitive types raises profit at every N ; and the optimal menu size lies in a bounded range. Simulations further reveal that the profit-maximizing complexity strictly exceeds the welfare-maximizing complexity, providing intuition for complexity regulation in markets with heterogeneous consumers. Our results connect classical screening theory to the literatures on consideration set formation, choice overload, and incentive complexity.

1 A Guide for the Reader

This section summarizes the paper's argument in ordinary language, with no equations, for readers who want the intuition first—or instead of—the formal derivations. Readers who prefer to start with the model can go directly to Section 2.

What the paper does. A firm has to decide how many versions of a product to sell—how many pricing tiers, subscription plans, or quality levels. Standard economic theory says: offer as many as you can profitably tell apart, because every additional tier lets you charge closer to what each customer is actually willing to pay. That logic assumes customers can always make sense of whatever menu they are shown, no matter how large. We ask what happens to that logic once some customers—not all, just some—find large menus hard to process. We build a model in which

the firm picks a single number, the number of tiers N , and then customers of varying cognitive ability respond to that menu as best they can. The paper asks: what number of tiers maximizes the firm’s profit once this is taken into account, and does it differ from the number that maximizes total social welfare?

The underlying idea, in one sentence. Adding another tier to a menu has two opposing effects: it lets the firm price-discriminate more finely among customers who can handle it (a good thing for profit), but it pushes some customers past their cognitive limit, causing them to walk away, skim only part of the menu, or pick something they would not have chosen with a clear head (a bad thing for profit, because a customer who leaves or picks badly is a customer who is not paying full price, or is not paying at all). At a small number of tiers, the first effect dominates and profit rises with complexity. Eventually, if the firm keeps adding tiers, almost every customer becomes overwhelmed, and profit is driven back down to zero. Since profit starts positive, and must eventually hit zero, and moves continuously between these two points as tiers are added one at a time, there has to be some number of tiers in between where profit is highest. That is the entire logical content of our headline result—nothing more exotic is going on.

Why we call it a “Laffer curve.” The original Laffer curve makes exactly this kind of argument about tax rates: revenue is zero at a 0% tax rate, revenue is zero at a 100% tax rate (nobody bothers to earn taxable income), and revenue is positive somewhere in the middle, so there must be a revenue-maximizing rate strictly between the two extremes. Nothing about that argument requires the revenue curve to rise smoothly and then fall smoothly with a single peak; it only requires an interior maximum to exist. We use the term in exactly this limited sense. We do *not* claim that our profit curve is single-peaked, that it has a unique maximum, or that it is concave in the number of tiers—we are unable to establish any of these stronger properties, and we say so explicitly wherever it matters (see the discussion after Theorem 12 below). The name refers to the existence of an interior optimum, not to the shape of the curve around it.

What is new, and how it fits with existing work. Two existing literatures come close to this question but stop short of asking it directly. One literature treats simplicity as a hard constraint on what a mechanism designer is *allowed* to do (for example, requiring that a mechanism be easy enough for consumers to understand at all). A second literature studies firms that deliberately exploit consumer confusion as a competitive weapon. Our firm does neither: it faces no simplicity requirement, and confusing customers is not in its interest, since a confused customer is disproportionately likely to walk away without buying anything. Instead, complexity is simply one more variable the firm chooses, trading off finer screening against behavioral breakage, in the same way it might choose a price or a quality level. That framing—complexity as an ordinary choice variable rather than a constraint or a weapon—is the paper’s main conceptual contribution. We view this as opening up a specific, tractable question rather than settling the general problem of mechanism design under bounded rationality: we restrict attention to a single, analytically convenient family

of menus (described in Section 4.2) and ask only how many tiers to use *within* that family, rather than solving for the unrestricted optimal menu when customers are boundedly rational. Section 3 places the paper relative to the specific literatures it draws on.

How to read the formal results: what is a destination and what is a stepping stone.

The paper contains many labeled definitions, assumptions, propositions, and theorems. Most of these are bookkeeping or intermediate steps needed to make the two main results watertight. Concretely:

- **Two results are the actual destinations.** Theorem 12 (interior existence, in the general model) and Theorem 16 (the sharper characterization available in a simplified special case) are what the paper is ultimately arguing for. Everything else supports one of these two theorems.
- **Assumption 11**, the “Interior Optimum Condition,” says in words: going from one tier to two tiers has to be worth it. Concretely, it asks that not too many customers—fewer than about 31% in the sufficient version we use—have so little cognitive capacity that they cannot even handle a two-tier menu. If almost everyone can cope with two tiers, expanding from one to two is a safe move for the firm, which is the first building block needed to get complexity off the ground at all.
- **Theorem 12** then says, in words: profit starts positive at one tier (part i); profit is driven to exactly zero once the menu is so large that every customer, even the most capable one in the population, gives up (part ii); and because a positive number can’t fall to zero without an interior point along the way at least as high as where it started, there must be some number of tiers that is at least as profitable as any other (part iii). Part (iv) is just the calculus-free version of “the profit-maximizing point is a local peak”: profit was still rising the step before it, and is no longer rising the step after it. We stress again what this theorem does *not* say: it does not say there is only one such best point, and it does not say profit rises smoothly to that point and falls smoothly afterward.
- **Theorem 16** asks what more can be said if we simplify the behavioral story so that a confused customer either walks away entirely or, if she doesn’t walk away, evaluates the menu just as carefully as anyone else (i.e., we turn off the “skims only part of the menu” and “chooses noisily” channels, keeping only exit). In this simplified world we can say quite a bit more precisely:
 - *Part (i), Condition R, in words:* a customer who is cognitively constrained never generates more expected profit, on average, than a customer facing the identical menu who is fully rational—and generates strictly less whenever there is a real chance she is constrained enough to walk away. This is intuitive once exit is the only friction: a constrained customer either leaves (zero profit to the firm) or stays and behaves rationally

(ordinary profit), so on average she can only do as well as or worse than a fully rational customer, never better.

- *Part (ii), in words:* total profit equals what the firm would earn from a fully rational population, scaled down by the fraction of customers who do not walk away. As the menu grows, the “fully rational” profit number would keep rising, but the fraction of customers who stick around keeps falling; the best menu size is wherever the product of a rising number and a falling number is largest.
 - *Part (iii), in words:* if we compare two markets, one where customers are, on average, more cognitively able than the other in a precise statistical sense (first-order stochastic dominance), the more able market is at least as profitable at *every* menu size, not just at the optimal one.
 - *Part (iv), in words:* the best menu size in this simplified model is never just one tier and never so large that everyone has already given up—it sits somewhere strictly in between, and we can write down an explicit number beyond which it is guaranteed not to lie.
- **Everything else**—the surplus bound (Proposition 19), the welfare propositions (Propositions 22–23), the Stackelberg formalities (Definition 7 and Proposition 8), and the profit decomposition (equation 7)—exists to make the two headline theorems precise and to set up the welfare comparison. None of these is intended to be a standalone economic punchline; they are the scaffolding.
 - **The welfare gap** ($N^W = 1 < N^* = 4$ in our baseline calibration) and the claim that the profit-maximizing profile looks single-peaked in practice are simulation findings, not theorems. We flag this distinction every time it is relevant, including in the discussion of policy implications in Section 7.

Bottom line. The formal apparatus exists to nail down one simple idea: once some customers can be overwhelmed by choice, a firm that ignores this fact and keeps adding tiers is leaving money on the table past some point, and a firm that accounts for it will stop short of full-blown screening. Exactly how far short, and whether that stopping point is also good for society, are questions we can answer sharply in a simplified special case and only illustrate by simulation in the general case.

2 Introduction

How many options should a monopolist offer? Classical screening theory answers unambiguously: as many as needed to separate consumer types fully (Mussa and Rosen 1978; Maskin and Riley 1984). With fully rational consumers, richer menus yield finer price discrimination and higher profit; the optimal menu is bounded only by the number of consumer types. This logic rests on one critical assumption—that consumers evaluate any menu effortlessly, regardless of its size.

That assumption is increasingly difficult to defend. Over two decades of empirical research document *choice overload*: confronted with large choice sets, consumers engage less, choose worse, and are more likely to abandon the decision altogether (Iyengar and Lepper 2000; see Chernev, Böckenholt, and Goodman 2015 for a meta-analysis). In financial markets, Huberman, Iyengar, and Jiang (2007) show that 401(k) participation falls as the number of fund options rises. In online retail, Natan (2025) estimates a structural model in which consumers consider only a fraction of available products and consideration probability declines with assortment size. In the laboratory, Abeler, Huffman, and Raymond (2025) demonstrate that complex incentive schemes distort effort provision relative to simpler schemes with identical expected payoffs—not because workers are irrational in general, but because the complex structure exceeds the cognitive budget needed to evaluate it fully. Yet no formal theory of monopoly pricing derives the profit-maximizing menu size when the designer accounts endogenously for the cognitive cost of complexity.

We develop such a theory. A monopolist faces a continuum of consumers who differ in both their valuation type θ (willingness to pay) and their cognitive type k (ability to evaluate complex menus). Given a parametric family of quality-price menus indexed by the number of tiers N , the monopolist commits to N to maximize expected profit. Consumers with cognitive type $k \geq N$ evaluate the full menu rationally; those with $k < N$ may exit with probability increasing in cognitive overload, evaluate only a randomly drawn subset of items, and choose with greater decision noise. This three-channel behavioral model—exit, restricted attention, and noise—captures in reduced form the qualitative effects that any model of cognitive resource exhaustion would predict.

The Complexity Laffer Curve. Our main result (Theorem 12) establishes that expected profit attains an interior maximum in N : for small N , the screening gain from adding tiers dominates; for sufficiently large N , behavioral losses dominate and profit eventually collapses to zero. We call this non-monotone pattern the **Complexity Laffer Curve**. The analogy to the Laffer Curve in public finance (Laffer 1981) is deliberate: just as raising tax rates beyond a point shrinks the tax base and reduces revenue, raising menu complexity beyond N^* erodes the base of effective participants and reduces monopoly profit. As with the original Laffer curve, the claim being made is only that an interior optimum exists because the function is positive somewhere in the interior and zero at the boundary; it is not a claim about the shape of the curve between those points (see the discussion in Section 1).

Two important caveats apply. First, Theorem 12 proves *existence* of an interior maximum—it does not prove the profit profile is unimodal or globally single-peaked; local optima may exist. Second, the monopolist here is not a government but a mechanism designer who internalizes the behavioral cost of complexity, which is why she does not push N to infinity.

Endogenous Complexity. We formalize the interaction as a Stackelberg commitment game: the monopolist (leader) commits to N and consumers (followers) best-respond under cognitive constraints. This situates our work between two adjacent literatures. The *complexity-as-constraint* literature (Li 2017; Pycia and Troyan 2023) treats simplicity as an implementability requirement.

The *complexity-as-weapon* literature (Spiegler 2006; Gabaix and Laibson 2006) studies firms that exploit consumer confusion. In our model, complexity is a *design variable*: the monopolist neither faces a simplicity constraint nor deliberately obfuscates; she optimizes over complexity itself to maximize profit.

Scope of the Contribution. We want to be precise about what kind of result this is. Solving for the fully unrestricted profit-maximizing selling mechanism when buyers are boundedly rational—optimizing simultaneously over the number of options, their qualities, their prices, and the allocation rule—is a substantially harder problem than the one we solve here, and we do not attempt it. Instead, we fix a specific, analytically tractable family of quality-price menus (Section 4.2) that already has the property economists would expect from classical screening theory—profit rises without bound in the number of tiers when consumers are rational—and we ask only how many tiers from that family the monopolist should deploy once some consumers are cognitively constrained. The payoff of this restriction is tractability: it lets us pin down an interior optimum with an explicit proof rather than a numerical exercise. The cost is scope: our theorems are statements about optimal complexity *within* this menu family, not about the unrestricted mechanism-design problem under bounded rationality. We view the paper’s contribution as identifying and cleanly answering a specific, previously unposed question—how much complexity should a designer choose when complexity is costly in a behavioral sense—rather than as a general theory of mechanism design with boundedly rational agents. We return to this scope limitation in Section 7.

Main Findings. *Existence (Theorem 12).* Under three explicit assumptions, expected profit attains an interior maximum $N^* \geq 2$. The sufficient condition (Assumption 11) requires that fewer than 31% of consumers have cognitive type below 2.

Exit-only characterization (Theorem 16). Under a tractable special case: (a) Condition R holds—constrained consumers yield weakly less profit than rational ones; (b) a FOSD shift toward higher cognitive types raises profit pointwise in N ; (c) $2 \leq N^* < N_0$ for an explicit upper bound.

Full-model FOSD (Conjecture 18). The pointwise profit ordering extends to the full model, as verified computationally. We discuss why the formal proof is analytically obstructed and leave it as an open problem.

Welfare gap (simulations). The profit-maximizing complexity $N^* = 4$ exceeds the welfare-maximizing complexity $N^W = 1$ in our baseline. This gap is a simulation finding, not a theorem; we discuss its economic interpretation and implications for complexity regulation.

Organization. Readers who want the intuition without the notation can consult Section 1 above, which walks through the argument, the role of each main assumption and theorem, and the distinction between proved results and simulation findings in plain language. Turning to the formal development: Section 3 surveys related literature. Section 4 presents the model. Section 5 develops

the main results. Section 6 presents simulation evidence. Section 7 discusses economic implications, limitations, and extensions. Section 8 concludes.

3 Related Literature

Classical Screening. Mussa and Rosen (1978) and Maskin and Riley (1984) characterize the optimal quality-price menu under private information with fully rational consumers. Profit is strictly increasing in the number of tiers; the monopolist optimally screens as finely as the type distribution warrants. Our Lemma 3 confirms this within our parametric family. The central departure is bounded rationality: beyond N^* , additional tiers reduce rather than raise profit.

Choice Overload and Large Assortments. Iyengar and Lepper (2000) document that consumers presented with a large jam display (24 varieties) were significantly less likely to purchase than those shown a small display (6 varieties), despite the larger set containing more potentially preferred options. The effect is real but heterogeneous across contexts (Chernev, Böckenholt, and Goodman 2015). In defined-contribution pension markets, Huberman, Iyengar, and Jiang (2007) find that participation falls by roughly two percentage points per ten additional fund options. Natan (2025) develops a structural demand model for large online assortments in which consideration probability declines with assortment size. Our model is a pricing-theoretic counterpart: we derive the profit-maximizing menu size rather than taking it as given.

Consideration Sets and Bounded Attention. Masatlioglu, Nakajima, and Ozbay (2012) provide an axiomatic characterization of attention in choice: from binary choice data, they identify which options “attract” attention and recover the consideration set correspondence alongside preferences. A key implication is that adding options can hurt consumers by crowding out attention from previously considered alternatives—a mechanism captured in our model through the uniform subset draw. Lleras, Masatlioglu, Nakajima, and Ozbay (2017) model how larger choice sets reduce the probability that any given option is considered, generating “when more is less” patterns in both choice frequencies and welfare. Our behavioral model is directly consistent with their framework: the consideration set size $m = \min(\lfloor k \rfloor, N - 1)$ in Assumption 5 falls relative to menu size as N grows, matching their key mechanism.

In the context of competitive markets, Eliaz and Spiegler (2011) model duopolistic competition when consumers form consideration sets through costly search. A key finding is that consideration set formation generates market power even among ex-ante symmetric firms. Our analysis is complementary: we study a monopolist who cannot influence attention probabilities but optimizes the menu size given the induced attention process.

Rational Inattention. Sims (2003) proposes that information processing capacity is scarce and costly, and that optimal information acquisition can be characterized by mutual information costs. Our behavioral model shares the spirit of rational inattention: consumers with lower k facing a

menu of size $N > k$ have less residual capacity, leading to exit, restricted consideration, and decision noise. The three channels correspond to the three ways an information-constrained agent can fall short of rational evaluation. We treat the behavioral model as a reduced form for tractability rather than deriving it from full entropy minimization.

Broader Behavioral Industrial Organization and Menu-Cost Literatures. Our paper also connects to two literatures we engage with only briefly in the main text. Behavioral industrial organization studies how firms and consumers interact when consumers depart from full rationality; DellaVigna (2009) surveys field evidence on present bias, limited attention, and other departures from standard preferences that shape firm strategy, and Ellison (2006) reviews how models of bounded rationality have been incorporated into industrial organization, including settings in which firms choose contract or price complexity in response to consumer confusion. Grubb (2015) surveys models of consumer overconfidence and biased beliefs in the specific context of nonlinear pricing and tariff design, a setting closely related to our menu family. On the rational-inattention side, Caplin and Dean (2015) and Matejka and McKay (2015) develop discrete-choice foundations for information-constrained decision-making, providing one possible microfoundation for the reduced-form noise channel in Assumption 5 that we do not pursue formally here. Finally, our use of the number of tiers as a costly choice variable echoes the older menu-cost literature in macroeconomics (Mankiw 1985), where firms trade off the benefit of adjusting prices against a fixed cost of doing so; our menu cost is behavioral rather than physical, but the underlying logic—that more adjustment or more complexity is not free and the optimizing agent internalizes that cost—is the same.

Complexity in Mechanism Design. Li (2017) introduces obviously strategy-proof (OSP) mechanisms and identifies which social choice functions are implementable at each simplicity level; Pycia and Troyan (2023) extend this to a richer theory. Both papers treat simplicity as an exogenous implementability constraint. Our question is orthogonal: within a fixed mechanism class, what level of complexity should the designer *choose*? Babaioff, Gonczarowski, and Nisan (2022) study how few menu items suffice to approximate optimal revenue with rational buyers; we study how many are too many with bounded rationality.

Complexity as a Strategic Tool. Spiegel (2006) and Gabaix and Laibson (2006) study competitive markets where firms strategically exploit consumer inattention. In those models, complexity is a weapon that benefits firms at consumers’ expense. Our monopolist is distinct: she wants consumers to participate and purchase, so excess complexity that triggers exit is directly harmful to her profit. This is precisely what generates the interior optimum.

Incentive Complexity and Field Evidence. Abeler, Huffman, and Raymond (2025) provide field and laboratory evidence that complex piece-rate schemes distort effort provision relative to equivalent simpler schemes. Workers under complex schemes provide roughly 10% less effort,

with the distortion concentrated among those with lower numeracy—consistent with our model’s prediction that behavioral distortions are worse for lower- k types.

4 Model

4.1 Environment

A monopolist sells a product of variable quality $q \geq 0$ to a unit mass of consumers. Each consumer has a private type (θ_i, k_i) : a *valuation type* $\theta_i \sim U[0, 1]$ and a *cognitive type* $k_i \sim G$ on $[1, \bar{k}]$, drawn independently. The valuation type determines willingness to pay: utility from option j at quality q_j and price t_j is $u = \theta q_j - t_j$. The cognitive type k parameterizes the ability to evaluate complex menus; higher k means greater capacity. The outside option yields $u = 0$. The monopolist’s cost is $C(q) = q^2/2$. The assumption $k \geq 1$ ensures every consumer can evaluate at least one item.

The monopolist observes G but not individual realizations (θ_i, k_i) .

Notation. Let $q_0 = t_0 = 0$ denote the outside option. Under indifference, the consumer selects the lowest-indexed option.

4.2 The Equal-Type-Partition Menu Family

We optimize menu size N within a parametric family $\{\mathcal{M}(N)\}_{N=1}^{\infty}$. This is a *restricted mechanism class*: we optimize the number of tiers within a specific analytically tractable family rather than solving for the fully unrestricted optimal mechanism under bounded rationality. The family is constructed so that, with rational consumers, profit $\Pi^{RAT}(N)$ is strictly increasing in N (Lemma 3), converging to the full-screening benchmark $1/12$. This makes it a natural vehicle for studying the tension between screening gains (which increase in N under rationality) and behavioral losses (which also increase in N).

Definition 1 (Equal-Type-Partition Menu). *Divide $[1/2, 1]$ into N equal-width intervals with cut-offs $\theta_j = 1/2 + j/(2N)$ for $j = 0, \dots, N$. For tier $j \in \{1, \dots, N\}$:*

$$q_j = \begin{cases} \max(0, 2\theta_{j-1} - 1) & \text{if } j < N, \\ (\theta_{N-1} + \theta_N)/2 & \text{if } j = N, \end{cases} \quad (1)$$

$$t_1 = \theta_0 q_1, \quad (2)$$

$$t_j = \theta_{j-1}(q_j - q_{j-1}) + t_{j-1} \quad \text{for } j = 2, \dots, N. \quad (3)$$

The pricing rule sets price differences so that each consumer type θ_{j-1} is indifferent between tier j and tier $j - 1$ —the standard binding downward incentive constraint in screening. Quality levels are proportional to the lower endpoint of the assigned type interval, ensuring incentive compatibility and individual rationality under rational choice.

Remark 2 (Dummy First Tier). *For $N \geq 2$, tier 1 has $q_1 = 0$ and $t_1 = 0$, payoff-equivalent to the outside option under rational choice. It is a bookkeeping device; active screening occurs across tiers $2, \dots, N$. For $N = 1$, the single tier has $q_1 = 3/4$ and $t_1 = 3/8$, serving types $\theta > 1/2$.*

Lemma 3 (Rational Profit).

$$\Pi^{RAT}(N) = \frac{16N^3 - 16N + 9}{192N^3} = \frac{1}{12} - \frac{1}{12N^2} + \frac{3}{64N^3}. \quad (4)$$

In particular, $\Pi^{RAT}(1) = 3/64$ and $\Pi^{FS} \equiv \lim_{N \rightarrow \infty} \Pi^{RAT}(N) = 1/12$. The sequence is strictly increasing:

$$\Pi^{RAT}(N+1) - \Pi^{RAT}(N) = \frac{32N^3 + 21N^2 - 11N - 9}{192N^3(N+1)^3} > 0 \quad \forall N \geq 1. \quad (5)$$

Proof. The closed form is derived in Appendix A. The numerator of (5), $p(N) = 32N^3 + 21N^2 - 11N - 9$, satisfies $p(1) = 33 > 0$ and $p'(N) = 96N^2 + 42N - 11 > 0$ for all $N \geq 1$. $\square$

Remark 4 (Rational Monotonicity and Its Limits). *Lemma 3 confirms the classical screening logic: with fully rational consumers, finer menus always raise profit. Adding a tier separates an additional pair of consumer types, generating incremental surplus that the monopolist can capture. Our Theorem 12 shows this monotonicity does not survive bounded rationality: profit attains an interior maximum rather than increasing without bound. Note that $\Pi^{FS} = 1/12$ is the limiting profit within this parametric family—not the unrestricted monopoly optimum under bounded rationality.*

4.3 Behavioral Model

The behavioral model specifies how cognitively constrained consumers respond to a menu that exceeds their processing capacity. Three channels operate sequentially:

- *Exit* (extensive margin): consumers facing an overwhelming menu may abandon the decision entirely (Iyengar and Lepper 2000).
- *Restricted consideration* (attention margin): those who engage evaluate only a subset of available options (Masatlioglu, Nakajima, and Ozbay 2012; Lleras et al. 2017; Eliaz and Spiegler 2011).
- *Decision noise* (quality margin): even conditional on evaluating a subset, larger menus are associated with higher noise in the selection process, consistent with rational inattention (Sims 2003).

Assumption 5 (Cognitive Cost Structure). *If $k \geq N$ (**sophisticated**): the consumer selects $\min \arg \max_{j \in \{0, \dots, N\}} (\theta q_j - t_j)$.*

*If $k < N$ (**constrained**), with $d = N - k > 0$:*

1. **Exit stage:** with probability $\gamma(d, k) = \min(1, \lambda d/(k+1))$, $\lambda > 0$, the consumer takes the outside option and exits. With probability $1 - \gamma(d, k)$, she proceeds.
2. **Evaluation stage** (conditional on not exiting): she evaluates $m = \min(\lfloor k \rfloor, N - 1)$ items drawn uniformly without replacement from $\{1, \dots, N\}$, forming consideration set S .
3. **Choice stage** (conditional on not exiting): she perceives $\hat{u}_j = \theta q_j - t_j + \varepsilon_j$ for each $j \in S \cup \{0\}$, where $\varepsilon_j \stackrel{iid}{\sim} \mathcal{N}(0, \eta^2 d^2 / N^2)$, $\eta > 0$, and $\varepsilon_0 = 0$. She selects $\min \arg \max_{j \in S \cup \{0\}} \hat{u}_j$.

On the Exit Probability. The form $\gamma(d, k) = \min(1, \lambda d/(k+1))$ increases in cognitive overload $d = N - k$ and decreases in capacity k : a barely-constrained consumer exits rarely; one whose capacity is severely exceeded exits with high probability. The parameter $\lambda > 0$ governs overall sensitivity. By construction, $\gamma = 0$ for sophisticated consumers ($d \leq 0$). The specific functional form is used in Theorem 12(ii) to derive the explicit bound N_0 ; the existence result in Theorem 12(iii) extends to any exit rule satisfying (a) sophistication when $k \geq N$, (b) existence of some N_0 beyond which all consumers exit, and (c) the sufficient condition in Assumption 11. The specific formula for N_0 does depend on the functional form of γ .

On the Noise Variance. The noise variance $\eta^2 d^2 / N^2$ scales with $(d/N)^2$: low when overload is small relative to menu size, high when a consumer with limited capacity faces a large menu. This is consistent with rational inattention: residual uncertainty grows when the information budget is spread more thinly across more items.

Assumption 6 (Regularity). G has support $[1, \bar{k}]$, $\bar{k} > 2$, continuous density $g > 0$ on $(1, \bar{k})$, $G(1) = 0$, and finite mean and variance.

4.4 Game-Theoretic Formulation

Definition 7 (Complexity Commitment Game). *The game has two stages:*

1. **Leader (monopolist):** commits to complexity $N \in \mathbb{N}$ and offers $\mathcal{M}(N)$.
2. **Followers (consumers):** each consumer (θ_i, k_i) observes $\mathcal{M}(N)$ and selects an option according to Assumption 5.

The monopolist's payoff is $\Pi(N)$, and the leader's problem is $\max_{N \in \mathbb{N}} \Pi(N)$.

The Stackelberg formulation is natural: the monopolist designs the menu before consumers arrive (e.g., prints a product catalog or configures a pricing page), so the commitment to N is observable and credible. The monopolist anticipates the distribution G but cannot condition on individual realizations.

Proposition 8 (Stackelberg Structure). *For each N , Assumption 5 defines a stochastic choice rule $\beta_N : [0, 1] \times [1, \bar{k}] \rightarrow \Delta(\{0, \dots, N\})$ mapping each type (θ, k) to a probability distribution over*

options, with randomness induced by the exit, subset, and noise draws. The expected payoff $\Pi(N) = E_{(\theta,k)}[\pi(\theta,k;N)]$ is a well-defined real number for each N . A Stackelberg-optimal commitment $N^* \in \arg \max_N \Pi(N)$ exists and is interior under the conditions of Theorem 12.

Proof. For each (θ, k, N) , the consumer executes a finite sequence of measurable randomizations. The composite rule $\beta_N(\theta, k)$ is a well-defined probability measure on $\{0, \dots, N\}$, and $\Pi(N)$ is a finite real number for each N . Interior existence follows from Theorem 12. $\square$

Definition 9 (Choice and Profit). *Let j^* denote the option selected by the consumer. The exit draw, subset draw, and noise draws are all internal to the consumer's decision process; we write $E[\cdot \mid k]$ for the expectation over $\theta \sim U[0, 1]$ and over all random components of the behavioral process (exit, subset draw, and noise draws) for a consumer of cognitive type k .*

Expected utility of type k : $U(k, N) = E[u^* \mid k]$, where $u^* = \theta q_{j^*} - t_{j^*}$.

Per-consumer expected profit: $\pi(\theta, k; N) = E[(t_{j^*} - C(q_{j^*})) \cdot \mathbf{1}(j^* \geq 1) \mid \theta, k]$, where the expectation integrates over all random components of the behavioral process.

Rational benchmark: $j^{RAT}(\theta, N) = \min \arg \max_{j \in \{0, \dots, N\}} (\theta q_j - t_j)$, $u^{RAT}(\theta, N) = \theta q_{j^{RAT}} - t_{j^{RAT}}$, $U^{RAT}(N) = \int_0^1 u^{RAT}(\theta, N) d\theta$, $CS^{RAT}(N) = U^{RAT}(N)$, and $W^{RAT}(N) = \Pi^{RAT}(N) + CS^{RAT}(N)$.

Total expected profit:

$$\Pi(N) = \int_1^{\bar{k}} \int_0^1 \pi(\theta, k; N) d\theta dG(k). \quad (6)$$

A key decomposition separates sophisticated from constrained consumers. Since all consumers with $k \geq N$ are sophisticated, they contribute $\Pi^{RAT}(N)$ per unit mass. Letting $\Pi^{CON}(N) \equiv E[\pi(\theta, k; N) \mid k < N]$ denote expected per-consumer profit from the constrained population:

$$\Pi(N) = [1 - G(N)] \Pi^{RAT}(N) + G(N) \Pi^{CON}(N). \quad (7)$$

Remark 10 (Direction of Behavioral Distortion). *We do not impose $\Pi^{CON}(N) \leq \Pi^{RAT}(N)$ as an assumption. This ordering is proved for the exit-only model (Theorem 16(i)) but treated as an empirical regularity for the full model (Remark 17). Intuitively, constrained consumers exit more and evaluate less carefully, reducing profit. However, decision noise could push a constrained consumer toward a higher-priced tier, raising profit on a given realization; whether the average effect is negative depends on parameterization.*

5 Main Results

5.1 Existence of an Interior Maximum

Assumption 11 (Interior Optimum Condition).

$$[1 - G(2)][\Pi^{RAT}(2) - \Pi^{RAT}(1)] > G(2)[\Pi^{RAT}(1) - \Pi^{CON}(2)].$$

Since $\Pi^{RAT}(2) - \Pi^{RAT}(1) = 11/512$ and $\Pi^{CON}(2) \geq 0$, a sufficient condition is $G(2) < 11/35 \approx 0.314$.

Economic Interpretation. Assumption 11 ensures it is profitable to expand from $N = 1$ to $N = 2$. The left-hand side is the *screening gain*: the fraction of consumers sophisticated at $N = 2$ (those with $k \geq 2$) times the additional rational profit from a second tier. The right-hand side is the *behavioral cost*: the fraction of consumers constrained at $N = 2$ (those with $k < 2$) times their profit shortfall relative to the rational benchmark. The condition holds automatically when fewer than 31% of consumers have cognitive type below 2.

Theorem 12 (Interior Profit Maximum). *Under Assumptions 5–11:*

- (i) (Baseline) $\Pi(1) = \Pi^{RAT}(1) = 3/64$.
- (ii) (Profit collapse) *There exists $N_0 = \lceil \bar{k} + (\bar{k} + 1)/\lambda \rceil$ such that $\Pi(N) = 0$ for all $N \geq N_0$. This bound depends on the specific functional form $\gamma(d, k) = \min(1, \lambda d/(k + 1))$ in Assumption 5.*
- (iii) (Interior existence) *There exists $N^* \in \{2, \dots, N_0 - 1\}$ with $\Pi(N^*) \geq \Pi(N)$ for all N . In particular, $\Pi(N^*) > \Pi(1)$. The profit profile need not be unimodal; the theorem establishes existence of a maximizer, not its uniqueness.*
- (iv) (Discrete optimality) *At any maximizer N^* : $\Delta\Pi(N^* - 1) \geq 0 \geq \Delta\Pi(N^*)$, where $\Delta\Pi(N) \equiv \Pi(N + 1) - \Pi(N)$.*

In Plain Words. Profit equals the fully rational benchmark when the menu is trivially small (part i); profit is driven all the way to zero once the menu is so large that even the most cognitively able consumer in the population gives up (part ii); since a quantity that starts positive cannot reach zero without passing through, or peaking before, an interior point, some menu size at least as good as any other must exist (part iii); and at that best menu size, the last tier added was worth adding but the next one would not have been (part iv). No claim is made that this best menu size is unique or that profit moves smoothly up to it and back down.

Proof. (i) At $N = 1$, all consumers satisfy $k \geq 1 = N$, so all are sophisticated and the behavioral model coincides with rationality.

(ii) The exit probability $\gamma(N - k, k) = 1$ when $N \geq k + (k + 1)/\lambda$. The right-hand side is increasing in k (since $\partial/\partial k[k + (k + 1)/\lambda] = 1 + 1/\lambda > 0$), so the maximum over $k \in [1, \bar{k}]$ is attained at $k = \bar{k}$. For $N \geq \lceil \bar{k} + (\bar{k} + 1)/\lambda \rceil = N_0$, every consumer exits regardless of type, so $\Pi(N) = 0$.

(iii) From (7):

$$\Pi(2) - \Pi(1) = [1 - G(2)][\Pi^{RAT}(2) - \Pi^{RAT}(1)] - G(2)[\Pi^{RAT}(1) - \Pi^{CON}(2)] > 0$$

by Assumption 11, so $\Pi(2) > \Pi(1) > 0$. Since $\Pi(N) = 0$ for $N \geq N_0$ and the feasible set $\{1, \dots, N_0 - 1\}$ is finite, a maximizer N^* exists over this set. Because $\Pi(2) > \Pi(1) \geq \Pi(N_0) = 0$, the maximizer satisfies $N^* \geq 2$.

(iv) Immediate from the definition of a discrete maximizer. $\square$

Remark 13 (Hump Shape vs. Interior Maximum). *Theorem 12 establishes that profit attains an interior maximum—it is positive at $N^* \geq 2$ and collapses to zero for large N —without asserting that the profile is monotonically increasing to the left of N^* and monotonically decreasing to the right. A globally hump-shaped (unimodal) profile would require additional structure, such as log-concavity of $\Pi(N)$ in N , which we do not impose. Simulations in Section 6 are consistent with unimodality in the parameterizations tested, but this is an empirical regularity rather than a theorem. We are unable to establish concavity, log-concavity, or uniqueness of the maximizer under the general conditions of Theorem 12, and we do not regard this as a gap that undermines the result: the economically meaningful claim—that complexity has a strategic ceiling, and profit does not rise forever—requires only existence of an interior optimum, exactly as the public-finance Laffer curve requires only that revenue be zero at both 0% and 100% tax rates and positive somewhere between them, not that the revenue function be single-peaked or concave.*

Marginal Analysis. Expanding $\Delta\Pi(N) = \Pi(N+1) - \Pi(N)$ using (7) with $\delta_G(N) = G(N+1) - G(N)$ (derivation in Appendix B):

$$\begin{aligned} \Delta\Pi(N) = & \underbrace{[1 - G(N)][\Pi^{RAT}(N+1) - \Pi^{RAT}(N)]}_{\text{screening gain}} \\ & - \underbrace{\delta_G(N)[\Pi^{RAT}(N+1) - \Pi^{CON}(N+1)]}_{\text{threshold cost}} \\ & + \underbrace{G(N)[\Pi^{CON}(N+1) - \Pi^{CON}(N)]}_{\text{effect on already-constrained consumers}}. \end{aligned} \quad (8)$$

The *screening gain* is always positive by Lemma 3. The *threshold cost* captures the profit loss from the mass $\delta_G(N)$ of consumers who become constrained when the menu expands from N to $N+1$; this term is non-negative when Condition R holds (which is proved for the exit-only model but not for the full model in general). The third term can be positive or negative. At the optimal N^* , the sum of these three terms changes sign: marginal gains no longer outweigh marginal costs.

5.2 Characterization under the Exit-Only Model

We obtain sharper analytical results under a tractable special case.

Definition 14 (Exit-Only Model). *Constrained consumers exit with probability $\gamma(d, k) = \min(1, \lambda d/(k+1))$; conditional on not exiting, they evaluate all N items rationally (no noise, no subset restriction).*

This model isolates the extensive margin of participation from the intensive margin of evaluation quality. Its economic logic is transparent: a constrained consumer either fails to engage entirely or, if she engages, makes fully rational decisions.

Definition 15 (Condition R). *Condition R holds if $\Pi^{CON}(N) \leq \Pi^{RAT}(N)$ for all $N \geq 1$.*

Theorem 16 (Exit-Only Characterization). *Under the exit-only model (Definition 14) and Assumptions 6–11:*

- (i) (Condition R) $\Pi^{CON}(N) \leq \Pi^{RAT}(N)$ for all N , with strict inequality whenever G assigns positive mass to $(1, N)$.
- (ii) (Profit formula) $\Pi(N) = [1 - \Gamma(N)] \cdot \Pi^{RAT}(N)$, where

$$\Gamma(N) = \int_1^{\bar{k}} \gamma(N-k, k) \mathbf{1}(k < N) dG(k)$$

satisfies $\Gamma(1) = 0$ and $\Gamma(N) = 1$ for all $N \geq N_0$.

- (iii) (FOSD monotonicity) If G_2 first-order stochastically dominates G_1 , then $\Pi(N; G_2) \geq \Pi(N; G_1)$ for every N .
- (iv) (Bounded optimizer) $2 \leq N^* < N_0$.

In Plain Words. When the only behavioral friction is exit, four things can be said cleanly. First, a constrained consumer is never more valuable to the firm, on average, than a fully rational one would be at the same menu, and is strictly less valuable whenever there is any real chance she gives up (part i, Condition R). Second, total profit is simply the rational-benchmark profit multiplied by the share of consumers who stick around, so the best menu size balances a rising benchmark against a shrinking retention rate (part ii). Third, a market with a more cognitively capable population is at least as profitable at every possible menu size, not only at its own optimum (part iii). Fourth, the profit-maximizing menu size is never the smallest possible menu and never so large that retention has collapsed to nothing—it is somewhere strictly in between, within an explicit, computable range (part iv).

Proof. (i) In the exit-only model, a constrained consumer of type $k < N$ exits with probability $\gamma(N-k, k) \in (0, 1]$ and, conditional on not exiting, chooses rationally. Her expected profit contribution is $(1 - \gamma(N-k, k)) \cdot \Pi^{RAT}(N)$. Averaging over the constrained population:

$$\Pi^{CON}(N) = E[(1 - \gamma(N-k, k)) \mid k < N] \cdot \Pi^{RAT}(N) \leq \Pi^{RAT}(N),$$

with strict inequality when $\gamma > 0$ on a set of positive G -measure in $(1, N)$, which holds under Assumption 6 whenever $N > 1$.

(ii) Substituting into (7):

$$\begin{aligned}\Pi(N) &= [1 - G(N)]\Pi^{RAT}(N) + G(N) \cdot E[1 - \gamma \mid k < N] \cdot \Pi^{RAT}(N) \\ &= [1 - \Gamma(N)]\Pi^{RAT}(N).\end{aligned}$$

At $N = 1$: the indicator $\mathbf{1}(k < 1) = 0$ for all $k \in [1, \bar{k}]$, so $\Gamma(1) = 0$. For $N \geq N_0$: $\gamma = 1$ for all $k \in [1, \bar{k}]$, so $\Gamma(N) = 1$.

(iii) Define $h(k, N) = \gamma(N - k, k) \mathbf{1}(k < N)$. On $k < N$: $h(k, N) = \min(1, \lambda(N - k)/(k + 1))$. The unsaturated branch has derivative $-\lambda(N + 1)/(k + 1)^2 < 0$ in k , and truncation at 1 preserves weak monotonicity. On $k \geq N$: $h = 0$. As $k \uparrow N$, the unsaturated branch tends to zero continuously, so h is weakly decreasing on $[1, \bar{k}]$. Under FOSD, $E_{G_2}[h(\cdot, N)] \leq E_{G_1}[h(\cdot, N)]$, giving $\Gamma(N; G_2) \leq \Gamma(N; G_1)$ and hence $\Pi(N; G_2) \geq \Pi(N; G_1)$.

(iv) From Theorem 12(ii)–(iii). □

Economic Interpretation. Part (i) (Condition R) says that cognitive constraints work against the monopolist in the exit-only model. This is transparent: a constrained consumer either exits (zero profit) or chooses rationally (rational profit), so her expected contribution is a fraction of rational profit, strictly below it.

Part (ii) gives an elegant formula: profit equals rational profit scaled by the participation-weighted non-exit rate $1 - \Gamma(N)$. As N increases, $\Pi^{RAT}(N)$ rises but $1 - \Gamma(N)$ falls. The interior maximum is where the product of these two opposing forces peaks.

Part (iii) has a natural economic reading: a more cognitively able population is more profitable at every menu size because behavioral losses are smaller. Consequently, monopolists serving more cognitively able markets can profitably offer more tiers: the profit-maximizing N^* is weakly larger under G_2 than G_1 (though uniqueness of N^* is not established, so the ordering of N^* values is guaranteed only when the argmax is unique).

Remark 17 (Open Problem: Full-Model Condition R). *Whether $\Pi^{CON}(N) \leq \Pi^{RAT}(N)$ holds under the full behavioral model (with noise and subset restrictions) is an open problem. The difficulty is that decision noise ε_j can push a constrained consumer toward a higher-priced tier than she would rationally choose, potentially raising profit on that realization. A formal proof would require showing that the expected profit gain from noise is dominated by the expected loss from exit and subset restriction. This holds in all parameterizations we tested but remains analytically unresolved.*

5.3 Comparative Statics under the Full Model

Theorem 16(iii) proves the pointwise FOSD profit ordering for the exit-only model. Extending it to the full model requires establishing that the noise and subset channels also preserve FOSD monotonicity. The obstacle is that the noise variance $\eta^2(N - k)^2/N^2$ depends on k , so a shift in G simultaneously changes both the composition of the constrained population and the per-consumer

noise level. Controlling for both effects simultaneously is technically demanding and we leave the formal proof for future work.

Conjecture 18 (FOSD Ordering, Full Model). *Under the full behavioral model (Assumption 5), if G_2 first-order stochastically dominates G_1 , then $\Pi(N; G_2) \geq \Pi(N; G_1)$ for every N . It follows that $\Pi^*(G_2) \equiv \max_N \Pi(N; G_2) \geq \max_N \Pi(N; G_1) \equiv \Pi^*(G_1)$. If moreover the maximizers $N^*(G_1)$ and $N^*(G_2)$ are unique, then $N^*(G_2) \geq N^*(G_1)$; without uniqueness, this ordering of maximizers need not hold in general.*

The pointwise ordering is verified computationally across three FOSD-ordered distributions in Section 6 (Table 2). The exit-only proof (Theorem 16(iii)) establishes the mechanism and provides the qualitative intuition.

5.4 Surplus Structure

Proposition 19 (Surplus Bound). *(i) For $k \geq N$: $U(k, N) = U^{RAT}(N)$. (ii) For $1 \leq k < N$: $U(k, N) \leq U^{RAT}(N)$, with strict inequality when the constraint changes the selected option with positive probability on a positive- θ -measure set.*

Proof. A constrained consumer either exits ($u^* = 0 \leq u^{RAT}(\theta, N)$ since $u^{RAT}(\theta, N) \geq 0$ by individual rationality) or selects from $S \cup \{0\} \subseteq \{0, \dots, N\}$, yielding realized utility $u^* = \theta q_{j^*} - t_{j^*} \leq \max_{j \in S \cup \{0\}} (\theta q_j - t_j) \leq \max_{j \in \{0, \dots, N\}} (\theta q_j - t_j) = u^{RAT}(\theta, N)$ by set inclusion. Taking expectations gives part (ii). Part (i) holds by the definition of sophisticated consumers. $\square$

Remark 20 (Note on the Bound). *The bound $u^* \leq u^{RAT}$ holds even when noise is present: the noise ε_j affects which j^* is selected, but the realized utility of the selected option is the deterministic quantity $\theta q_{j^*} - t_{j^*}$, which is bounded above by the maximum deterministic utility over any superset.*

Interpretation. Cognitive constraints are unambiguously bad for consumers: every consumer is weakly worse off than under rationality, and there is no compensating distortion benefiting the constrained. This makes the welfare analysis straightforward.

5.5 Welfare Analysis

Definition 21 (Total Welfare). $W(N) = \Pi(N) + CS(N)$, where $CS(N) = \int_1^{\bar{k}} U(k, N) dG(k)$.

Proposition 22 (Consumer-Surplus Bound). $CS(N) \leq CS^{RAT}(N)$ for all N , with equality at $N = 1$.

Proof. Integrate Proposition 19(ii) over G : $CS(N) = \int U(k, N) dG(k) \leq \int U^{RAT}(N) dG(k) = U^{RAT}(N) = CS^{RAT}(N)$. Equality at $N = 1$ holds because all consumers are sophisticated. $\square$

Proposition 23 (Welfare at $N = 1$). $W(1) = 9/64$, decomposed as $\Pi(1) = 3/64$ and $CS(1) = 3/32$.

Proof. At $N = 1$, all consumers are sophisticated. A consumer of type θ purchases if and only if $\theta > 1/2$, yielding utility $\theta \cdot (3/4) - 3/8 = 3(\theta - 1/2)/4$.

$$CS(1) = \int_{1/2}^1 \left(\theta \cdot \frac{3}{4} - \frac{3}{8} \right) d\theta = \left[\frac{3\theta^2}{8} - \frac{3\theta}{8} \right]_{1/2}^1 = \frac{3}{8} - \frac{3}{8} - \frac{3}{32} + \frac{3}{16} = \frac{3}{32}.$$

Adding $\Pi(1) = 3/64$: $W(1) = 6/64 + 3/64 = 9/64$. □

Remark 24 (Welfare Optimum: Analytical vs. Simulation Result). *Proposition 23 establishes only that $W(1) = 9/64$. It does not prove that $N^W = 1$ is the welfare-maximizing menu size. The claim that $N^W = 1 < N^* = 4$ is a simulation finding, reported in Table 5 and the surrounding discussion. Whether the welfare optimum is always $N^W = 1$ depends on the parameterization; in our baseline it is, but this need not hold in general.*

The intuition for the welfare gap—when it exists—is as follows. As N increases, the monopolist gains from finer price discrimination, but also creates behavioral deadweight loss: exit produces zero surplus on both sides, and distorted choices reduce consumer welfare without generating offsetting profit. The monopolist partially internalizes the participation cost (exit hurts her profit) but not the full consumer surplus cost of distorted choices. This wedge between private and social incentives can generate a welfare gap $N^W < N^$ in parameterizations where the distortion is sufficiently large.*

Remark 25 (Welfare and Rational Benchmark). *The bound $CS(N) \leq CS^{RAT}(N)$ does not imply $W(N) \leq W^{RAT}(N)$: profit and consumer surplus move in opposite directions across menus, and the comparison across rational and behavioral environments is non-trivial. We treat $W(N)$ comparisons across menu sizes as empirical questions answered by simulation.*

Definition 26 (Price of Cognitive Heterogeneity). $PoCH(G) \equiv 1 - \Pi^*(G)/\Pi^{FS}$, where $\Pi^*(G) = \max_N \Pi(N; G)$ and $\Pi^{FS} = 1/12$. *PoCH measures the fraction of full-screening profit lost to cognitive heterogeneity. We do not prove $PoCH(G) \geq 0$ in general; it is positive in all parameterizations tested.*

6 Simulation Evidence

Parameters. $\theta \sim U[0, 1]$; $C(q) = q^2/2$; $k = 1 + \text{Beta}(\alpha_G, \beta_G) \cdot 11$, so $k \in [1, 12]$, $\bar{k} = 12$; $\lambda = 0.30$; $\eta = 0.12$; $n = 15,000$ consumers per run; 10 replications (seeds $42 + r \cdot 1000 + N$, $r = 1, \dots, 10$). The simulation algorithm is in Appendix C.

Definitions. *Participation*: fraction with $j^* \geq 1$ and $q_{j^*} > 0$. *Efficiency*: fraction of participants whose choice matches the rational benchmark $j^{RAT}(\theta, N)$. *Surplus groups*: ex-ante median split of G .

We report five experiments: (1) profit and participation as a function of N ; (2) FOSD comparative statics; (3) the effect of cognitive dispersion at fixed mean; (4) surplus inequality by cognitive type; and (5) welfare decomposition.

6.1 Experiment 1: The Profit Profile

Table 1 reports profit, participation, and efficiency for selected $N \in \{1, \dots, 25\}$, with the full grid searched to find $N^* = 4$.

Profit rises from $\Pi(1) = 0.047$ (matching $3/64 \approx 0.047$) to $\Pi(4) = 0.074$, then falls. The profile is consistent with an interior maximum at $N^* = 4$; it appears approximately single-peaked in this parameterization, though Theorem 12 does not guarantee unimodality in general (Remark 13).

Efficiency is the leading indicator of profit erosion: it falls from 100% at $N = 1$ to 82% at $N = 4$ and 17% at $N = 10$, well before participation collapses (which occurs mainly for $N \geq 10$). This suggests that the intensive-margin channel (noise, subset restriction) dominates at moderate complexity, while the extensive-margin channel (exit) dominates at high complexity.

Table 1: Profit, Participation, and Efficiency as a Function of Menu Size

N	Profit	SE	Particip.	Efficiency
1	0.047	0.000	0.506	1.000
2	0.068	0.000	0.492	0.981
3	0.072	0.000	0.490	0.922
4	0.074	0.000	0.488	0.822
5	0.072	0.000	0.489	0.684
7	0.068	0.000	0.494	0.426
10	0.060	0.000	0.456	0.169
15	0.043	0.000	0.325	0.078
21	0.025	0.000	0.189	0.047
25	0.016	0.000	0.122	0.044

Notes: Baseline: $k = 1 + \text{Beta}(2, 3) \cdot 11$ (mean $\mu_k = 5.4$, variance $\sigma_k^2 = 4.84$). Bold row: $N^* = 4$. Standard errors are of order 10^{-4} ; they appear as 0.000 at three decimal places. The profit profile appears approximately single-peaked; unimodality is not proved in general.

6.2 Experiment 2: FOSD Shift in Cognitive Ability

Table 2 compares three populations ordered by FOSD (verified on a grid of 1,000 points), with mean cognitive type ranging from 3.5 (low) to 9.5 (high). Both the optimal menu size N^* and maximum profit Π^* increase monotonically with cognitive ability, consistent with Theorem 16(iii) and Conjecture 18. N^* increases from 3 (low-ability population) to 8 (high-ability population), illustrating that the optimal menu complexity depends substantially on the cognitive composition of the market. The ordering of N^* values is well-defined here because each profit profile has a

unique maximum in these parameterizations; in cases of non-unique argmax, only the ordering of Π^* is guaranteed by the FOSD result.

Table 2: FOSD Shift in Cognitive Ability

Population	Distribution	μ_k	σ_k^2	N^*	Π^*
Low	$1 + \text{Beta}(1.5, 5.0) \cdot 11$	3.5	2.9	3	0.065
Medium	$1 + \text{Beta}(3.0, 3.0) \cdot 11$	6.5	4.3	6	0.077
High	$1 + \text{Beta}(5.0, 1.5) \cdot 11$	9.5	2.9	8	0.082

Notes: FOSD ordering verified on a 1,000-point grid. All other parameters at baseline values. N^* ordering is well-defined here because each profit profile has a unique maximum in these parameterizations.

6.3 Experiment 3: Cognitive Dispersion at Fixed Mean

Table 3 holds the mean cognitive type fixed at $\mu_k = 6.5$ and varies dispersion. As dispersion increases, PoCH rises from 5.1% to 15.6% and N^* falls from 6 to 4. Greater dispersion implies a longer left tail of low-ability consumers who are severely constrained by complex menus, limiting the profitability of fine screening. This experiment varies the *spread* of the distribution at fixed mean; it is distinct from the FOSD shift in Experiment 2 and does not directly bear on Conjecture 18.

Table 3: Cognitive Dispersion at Fixed Mean ($\mu_k = 6.5$)

Population	Distribution	$\text{Var}(k)$	N^*	Π^*	PoCH
Homogeneous	$1 + \text{Beta}(8, 8) \cdot 11$	1.8	6	0.079	5.1%
Moderate	$1 + \text{Beta}(3, 3) \cdot 11$	4.3	6	0.077	7.6%
Heterogeneous	$1 + \text{Beta}(0.8, 0.8) \cdot 11$	11.6	4	0.070	15.6%

Notes: Mean $k = 6.5$ in all rows (held fixed by construction). $\text{Var}(k)$ denotes the variance of the cognitive type distribution. $\text{PoCH} = 1 - \Pi^* / \Pi^{FS}$, $\Pi^{FS} = 1/12$. This experiment varies dispersion at fixed mean and does not impose FOSD across rows.

6.4 Experiment 4: Surplus Inequality by Cognitive Type

Table 4 reports expected consumer surplus for high- k and low- k consumers (median split of G) across menu sizes. At $N = 1$, both groups receive identical surplus—consistent with Proposition 22, since all consumers are rational at $N = 1$. For $N \geq 2$, a surplus gap opens, peaking at $N = 7$ (gap of 0.016), then narrowing as participation collapses across both groups. Low-ability consumers bear a disproportionate share of the welfare cost of complexity: at the profit-maximizing menu ($N^* = 4$), their surplus is roughly 23% lower than that of high-ability consumers.

Table 4: Surplus Inequality by Cognitive Type

N	$U_{\text{high-}k}$	$U_{\text{low-}k}$	Gap	Particip.
1	0.094	0.094	0.000	0.502
2	0.029	0.027	0.002	0.495
4	0.030	0.023	0.006	0.493
7	0.032	0.015	0.016	0.491
11	0.023	0.009	0.015	0.435
21	0.010	0.001	0.009	0.190

Notes: Groups defined by ex-ante median split of G (baseline distribution $k = 1 + \text{Beta}(2, 3) \cdot 11$). Gap peaks at $N = 7$ and narrows at high N as participation collapses for both groups.

6.5 Experiment 5: Welfare Decomposition

Table 5 reports $\Pi(N)$, $CS(N)$, and $W(N)$ for selected menu sizes.

In this parameterization, total welfare is maximized at $N^W = 1$, with $W(1) = 0.141 = 9/64$ (as computed analytically in Proposition 23). Profit is maximized at $N^* = 4$, with $\Pi(4) = 0.074$. The welfare at the profit-optimal menu, $W(4) = 0.101$, is approximately 28% below $W(1)$.

We emphasize that $N^W = 1$ is a *simulation finding*, not a theorem. Proposition 23 establishes only that $W(1) = 9/64$; the claim that no other menu size achieves higher welfare is verified only within this parameterization. Across different cognitive type distributions or behavioral parameters, the welfare optimum need not be $N = 1$.

The sharp welfare drop from $N = 1$ to $N = 2$ —consumer surplus falls from 0.094 to 0.028, far outweighing the profit increase from 0.047 to 0.068—suggests that the primary welfare cost of complexity arises at the first step of expansion, where the behavioral deadweight loss is substantial and concentrated among consumers with $k \in (1, 2)$.

Table 5: Welfare Decomposition
by Menu Size

N	$\Pi(N)$	$CS(N)$	$W(N)$
1	0.047	0.094	0.141
2	0.068	0.028	0.096
4	0.074	0.027	0.101
7	0.068	0.024	0.092
11	0.056	0.016	0.072
21	0.025	0.006	0.031

Notes: Bold rows: welfare optimum in this parameterization ($N^W = 1$, simulation finding) and profit optimum ($N^* = 4$). $W(1) = 9/64 \approx 0.141$ (Proposition 23 computes this value analytically). $N^W = 1$ is verified by simulation, not proved analytically. Welfare gap: $W(1) - W(4) \approx 0.040$, approximately 28% of $W(1)$.

6.6 Experiment 6: Alternative Behavioral Calibration

Table 6 and the welfare finding of Experiment 5 both rest on a single calibration ($\lambda = 0.30$, $\eta = 0.12$). To gauge how much the qualitative pattern depends on this choice, we recompute N^* , N^W , Π^* , and the welfare gap under two alternative calibrations, holding the cognitive type distribution at its baseline $k = 1 + \text{Beta}(2, 3) \cdot 11$: a *high-friction* case with a larger exit sensitivity and more decision noise ($\lambda = 0.45$, $\eta = 0.20$), and a *low-friction* case with a smaller exit sensitivity and less decision noise ($\lambda = 0.15$, $\eta = 0.06$).

The qualitative pattern is unchanged across all three calibrations: the profit-maximizing menu size strictly exceeds the welfare-maximizing menu size, which remains $N^W = 1$ throughout. The *magnitude* of the gap, however, is calibration-sensitive: it ranges from 16% of $W(N^W)$ under low friction to 37% under high friction. We read this as evidence that the qualitative direction of the result—profit-maximizing complexity exceeds welfare-maximizing complexity—is not an artifact of one specific parameter choice, while the specific size of the welfare gap, and any policy magnitude derived from it, is. We return to this caveat when discussing policy implications in Section 7.

Table 6: Robustness to Behavioral-Parameter Calibration

Calibration	λ	η	N^*	N^W	Π^*	$W(N^W)$	Gap
Low friction	0.15	0.06	6	1	0.083	0.141	16%
Baseline	0.30	0.12	4	1	0.074	0.141	28%
High friction	0.45	0.20	3	1	0.061	0.141	37%

Notes: Cognitive type distribution held fixed at the baseline $k = 1 + \text{Beta}(2, 3) \cdot 11$ across rows; only the exit-sensitivity parameter λ and the noise-scale parameter η vary. Gap = $[W(N^W) - W(N^*)]/W(N^W)$. As in Experiment 5, $N^W = 1$ and the welfare gap are simulation findings, not theorems, in every row of this table.

6.7 Continuity of the Interior Optimum under Menu Perturbations

The result below is a mathematical continuity property of the equal-type-partition menu family, not an empirical robustness check in the sense of Experiment 6 above; we label it accordingly rather than as “robustness” to avoid overstating what it shows.

Proposition 27 (Continuity of the Interior Optimum). *Suppose $\{\mathcal{M}_\epsilon(N)\}$ satisfies: (a) $\Pi(N; \epsilon)$ is continuous in ϵ at $\epsilon = 0$ for each N ; (b) there exists $\tilde{N}_0 < \infty$ such that $\Pi(N; \epsilon) = 0$ for all $N \geq \tilde{N}_0$ and $|\epsilon| < \bar{\epsilon}$; (c) $\Pi(2; 0) > \Pi(1; 0)$. Then for sufficiently small $|\epsilon|$, an interior maximizer $N^*(\epsilon) \geq 2$ exists.*

Proof. By continuity and (c), $\Pi(2; \epsilon) > \Pi(1; \epsilon)$ for $|\epsilon|$ small. The feasible set $\{1, \dots, \tilde{N}_0 - 1\}$ is finite, so a maximizer $N^*(\epsilon) \geq 2$ exists. $\square$

This is a statement about mathematical continuity, not about economic robustness: it shows that the interior maximum does not depend on the exact functional form of the equal-type-partition family down to the last detail—any menu family that is a small, continuous perturbation of it and satisfies conditions (a)–(c) also exhibits an interior profit optimum. It does not speak to how sensitive the *location* of the optimum or the *size* of the welfare gap are to the behavioral parameters λ and η ; that empirical question is addressed directly by the alternative calibrations in Experiment 6 above.

7 Discussion

7.1 The Strategic Logic of Complexity

The central message is that complexity is a strategic design variable with an interior profit-maximizing level. The monopolist is not indifferent between 2 and 20 tiers: too few tiers sacrifice price discrimination; too many erode participation and evaluation quality. Theorem 12 proves this

gives rise to an interior maximum; the simulations show that in our baseline parameterization the resulting profit profile appears approximately single-peaked.

This trade-off is visible in practice. Streaming services (Netflix, Spotify, Apple TV+) typically offer two to four pricing tiers despite serving highly heterogeneous consumers; mobile carriers, Internet providers, and software-as-a-service firms routinely cap their menus at three to five plans. Our framework suggests that even absent operational or marketing reasons for simplicity, cognitive heterogeneity alone is sufficient to generate a finite optimal menu size—provided that the cognitive type distribution places meaningful mass near or below any given N .

7.2 Relationship to Mechanism Design

Li (2017) and Pycia and Troyan (2023) take simplicity as a constraint and ask what can be implemented within it. Our question is orthogonal: given a fixed mechanism class, what level of complexity should the designer commit to? This distinction parallels the difference between characterizing the set of incentive-compatible mechanisms and selecting the revenue-maximizing mechanism from that set: we add complexity as a third dimension of mechanism design.

The exit-only model (Theorem 16) provides the cleanest structural results because the only behavioral deviation is exit, making the economic logic transparent. The full model adds noise and subset restrictions, which are empirically important (efficiency falls sharply in Table 1) but analytically harder. The key open problem—proving Condition R and the FOSD ordering under the full model (Remark 17 and Conjecture 18)—is analogous to proving revenue equivalence under non-standard behavioral models: technically demanding because the mapping from behavior to profit involves non-linear mixtures.

7.3 Implications for Policy

The simulation finding $N^W = 1 < N^* = 4$ suggests that in markets resembling our baseline parameterization, the profit-maximizing monopolist offers menus that are more complex than is socially optimal. The mechanism is intuitive: the monopolist internalizes the participation cost of complexity (exit hurts her profit) but not the full consumer surplus cost (distorted choices reduce consumer welfare without fully reducing monopoly profit). When this wedge is large, a cap on the number of tiers—or a “complexity tax” per additional tier—could raise total welfare.

This intuition broadly parallels regulatory interventions already observed in practice. Financial product regulators (MiFID II in the EU, Reg BI in the US) impose disclosure requirements and product complexity restrictions. Health insurance marketplaces under the ACA standardize plan designs to reduce cognitive burden on consumers. Our model provides an economic framework for evaluating such interventions: the welfare gain from mandated simplification is the difference $W(N^W) - W(N^*)$, which in our baseline simulation equals roughly 28% of $W(N^W)$, and ranges from 16% to 37% across the alternative calibrations in Experiment 6. We caution that this magnitude is parameterization-specific; the appropriate complexity bound in a given market requires empirical estimation of the cognitive type distribution and behavioral parameters.

We want to be explicit that the policy discussion in this subsection is illustrative, not load-bearing. It rests on a small number of simulated calibrations of a single, restricted menu family, not on an estimated model fit to market data, and not on a theorem establishing that $N^W < N^*$ holds in general. Experiment 6 shows that the *direction* of the gap is not an artifact of one calibration, but the specific percentages above should be read as illustrating the mechanism, not as an estimate a regulator could act on directly. Any real complexity cap or “complexity tax” would require empirical work well beyond the scope of this paper.

7.4 Limitations and Directions for Future Work

Restricted Mechanism Class. We study a parametric family, not the fully unrestricted optimal mechanism under bounded rationality. Proposition 27 shows the interior maximum survives small, continuous perturbations of the menu family, but this is a mathematical continuity result, not evidence on how the optimum or the welfare gap move with the behavioral parameters (Experiment 6 addresses that empirical question directly). The unrestricted problem—optimizing over all (quality, price) menus under behavioral constraints—is substantially harder and left for future work.

Unimodality. Theorem 12 proves existence of an interior maximum without establishing unimodality (Remark 13). Proving that the profit profile is globally single-peaked would likely require additional monotonicity or log-concavity conditions on the profit function that we do not currently have.

Full-Model Analytical Results. Condition R and the FOSD conjecture remain open for the full behavioral model. Progress on these would likely require bounding expectations of nonlinear functions of multivariate normals under coupling arguments, which is technically demanding.

Competition and Dynamics. We study a static monopolist. In dynamic settings, firms update menus based on observed participation; in competitive settings, firms may strategically complicate menus to reduce switching or simplify to attract cognitively constrained consumers. Both extensions are important directions for future work.

8 Conclusion

We develop a theory of optimal menu complexity under bounded rationality. Within a parametric family of quality-price screening menus, the monopolist faces a fundamental trade-off: finer menus improve price discrimination but erode participation and decision quality among cognitively constrained consumers. The resolution—the Complexity Laffer Curve—is an interior profit-maximizing menu size N^* , whose existence is proved under explicit sufficient conditions (Theorem 12).

Under a tractable exit-only special case (Theorem 16), we prove that constrained consumers generate weakly lower profit than rational ones, that FOSD shifts in cognitive ability raise profit pointwise, and that the optimal menu size is bounded. These results have natural economic interpretations: greater cognitive ability supports finer menus; more heterogeneous markets call for simpler ones. Simulations reveal a welfare gap between the profit-optimal and welfare-optimal complexity in our baseline parameterization, providing intuition for complexity regulation.

We have been careful throughout to distinguish what is proved (Theorems 12 and 16), what is an open problem (Condition R and FOSD for the full model), and what is a simulation finding (unimodality of the profit profile, $N^W = 1$). We hope this framework contributes to the theory of mechanism design under bounded rationality and motivates further empirical and theoretical work on optimal complexity in pricing, insurance, and financial product markets.

9 Declarations

Availability of Data and Materials. This paper is based entirely on a theoretical model and simulated data. No external datasets were used. All simulation procedures, parameter choices, and replication steps are fully described in Appendix C.

Competing Interests. The author declares no competing interests.

Funding. This research received no external funding.

Author Contributions. The author is solely responsible for the conception, analysis, and writing of this manuscript.

Acknowledgements. The author is grateful to Dr. Ran Spiegler for helpful discussions and valuable feedback on an earlier version of this manuscript.

REFERENCES

- Abeler, Johannes, David Huffman, and Collin Raymond. 2025. “Incentive Complexity, Bounded Rationality, and Effort Provision.” *American Economic Review*, 115(12): 4404–37. doi:10.1257/aer.20230751
- Babaioff, Moshe, Yannai A. Gonczarowski, and Noam Nisan. 2022. “The Menu-Size Complexity of Revenue Approximation.” *Games and Economic Behavior*, 134: 281–307. doi:10.1016/j.geb.2021.03.001
- Caplin, Andrew, and Mark Dean. 2015. “Revealed Preference, Rational Inattention, and Costly Information Acquisition.” *American Economic Review*, 105(7): 2183–2203. doi:10.1257/aer.20140117
- Chernev, Alexander, Ulf Böckenholt, and Joseph Goodman. 2015. “Choice Overload: A Conceptual Review and Meta-Analysis.” *Journal of Consumer Psychology*, 25(2): 333–58. doi:10.1016/j.jcps.2014.08.001
- DellaVigna, Stefano. 2009. “Psychology and Economics: Evidence from the Field.” *Journal of Economic Literature*, 47(2): 315–72. doi:10.1257/jel.47.2.315
- Ellison, Glenn. 2006. “Bounded Rationality in Industrial Organization.” In *Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress*, Vol. 2, edited by Richard Blundell, Whitney K. Newey, and Torsten Persson, 142–74. Cambridge: Cambridge University Press.
- Eliaz, Kfir, and Ran Spiegler. 2011. “Consideration Sets and Competitive Marketing.” *Review of Economic Studies*, 78(1): 235–62. doi:10.1093/restud/rdq016
- Gabaix, Xavier, and David Laibson. 2006. “Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets.” *Quarterly Journal of Economics*, 121(2): 505–40. doi:10.1162/qjec.2006.121.2.505
- Grubb, Michael D. 2015. “Overconfident Consumers in the Marketplace.” *Journal of Economic Perspectives*, 29(4): 9–35. doi:10.1257/jep.29.4.9
- Huberman, Gur, Sheena S. Iyengar, and Wei Jiang. 2007. “Defined Contribution Pension Plans: Determinants of Participation and Contribution Rates.” *Journal of Financial Services Research*, 31(1): 1–32. doi:10.1007/s10693-006-0034-7
- Iyengar, Sheena S., and Mark R. Lepper. 2000. “When Choice Is Demotivating: Can One Desire Too Much of a Good Thing?” *Journal of Personality and Social Psychology*, 79(6): 995–1006. doi:10.1037/0022-3514.79.6.995
- Laffer, Arthur B. 1981. “Government Exactions and Revenue Deficiencies.” *Cato Journal*, 1(1): 1–21.
- Li, Shengwu. 2017. “Obviously Strategy-Proof Mechanisms.” *American Economic Review*, 107(11): 3257–87. doi:10.1257/aer.20160425
- Lleras, Juan S., Yusufcan Masatlioglu, Daisuke Nakajima, and Erkut Y. Ozbay. 2017. “When More Is Less: Limited Consideration.” *Journal of Economic Theory*, 170: 294–311. doi:10.1016/j.jet.2017.03.002
- Mankiw, N. Gregory. 1985. “Small Menu Costs and Large Business Cycles: A Macroeconomic Model of Monopoly.” *Quarterly Journal of Economics*, 100(2): 529–37. doi:10.2307/1885395
- Matejka, Filip, and Alisdair McKay. 2015. “Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model.” *American Economic Review*, 105(1): 272–98. doi:10.1257/aer.20130047

- Masatlioglu, Yusufcan, Daisuke Nakajima, and Erkut Y. Ozbay. 2012. “Revealed Attention.” *American Economic Review*, 102(5): 2183–2205. doi:10.1257/aer.102.5.2183
- Maskin, Eric, and John Riley. 1984. “Monopoly with Incomplete Information.” *RAND Journal of Economics*, 15(2): 171–96. JSTOR
- Mussa, Michael, and Sherwin Rosen. 1978. “Monopoly and Product Quality.” *Journal of Economic Theory*, 18(2): 301–17. doi:10.1016/0022-0531(78)90085-6
- Natan, Or. 2025. “Choice Frictions in Large Assortments.” *Marketing Science*, forthcoming.
- Pycia, Marek, and Peter Troyan. 2023. “A Theory of Simplicity in Games and Mechanism Design.” *Econometrica*, 91(4): 1495–1526. doi:10.3982/ECTA16310
- Sims, Christopher A. 2003. “Implications of Rational Inattention.” *Journal of Monetary Economics*, 50(3): 665–90. doi:10.1016/S0304-3932(03)00029-1
- Spiegler, Ran. 2006. “Competition over Agents with Boundedly Rational Expectations.” *Theoretical Economics*, 1(2): 207–31. doi:10.3982/TE15

A Rational Profit Derivation

We derive $\Pi^{RAT}(N) = (16N^3 - 16N + 9)/(192N^3)$.

Case $N = 1$. $q_1 = 3/4$, $t_1 = 3/8$. Per-consumer conditional profit: $t_1 - C(q_1) = 3/8 - 9/32 = 3/32$. Mass of purchasing consumers: $1/2$. Total: $\Pi^{RAT}(1) = (3/32)(1/2) = 3/64$.

Case $N \geq 2$. Cutoffs $\theta_j = 1/2 + j/(2N)$; each tier has type mass $1/(2N)$. Tier 1 is dummy (zero profit contribution). Below, each expression is the per-consumer conditional profit for a type in that tier; total contribution is obtained by multiplying by $1/(2N)$.

For $2 \leq j \leq N-1$: $q_j = (j-1)/N$, $t_j = (j-1)(2N+j)/(4N^2)$. Per-consumer conditional profit: $(j-1)(2N-j+2)/(4N^2)$.

Top tier $j = N$: $q_N = (4N-1)/(4N)$, $t_N = 3(2N^2-1)/(8N^2)$. Per-consumer conditional profit: $(8N^2+8N-13)/(32N^2)$.

Summing.

$$\Pi^{RAT}(N) = \frac{1}{2N} \left[\sum_{j=2}^{N-1} \frac{(j-1)(2N-j+2)}{4N^2} + \frac{8N^2+8N-13}{32N^2} \right].$$

Setting $m = j-1$: $\sum_{m=1}^{N-2} m(2N-m+1) = (N-2)(N-1)(2N+3)/3$. Combining (and verified by direct computation for $N = 2, 3, 4, 5$) yields $(16N^3 - 16N + 9)/(192N^3)$.

Verification: $\Pi^{RAT}(1) = 3/64$; $\Pi^{RAT}(2) = 35/512$; $\Pi^{RAT}(3) = 131/1728$; $\Pi^{RAT}(4) = 323/4096$; $\Pi^{RAT}(5) = 643/8000$. Transfer bound: $t_N < 3/4$ for all N .

B Discrete Optimality Condition

Starting from (7) and writing $\delta_G(N) = G(N+1) - G(N)$, substitute $1 - G(N+1) = [1 - G(N)] - \delta_G(N)$ and collect terms:

$$\begin{aligned} \Pi(N+1) - \Pi(N) &= [1 - G(N+1)]\Pi^{RAT}(N+1) + G(N+1)\Pi^{CON}(N+1) \\ &\quad - [1 - G(N)]\Pi^{RAT}(N) - G(N)\Pi^{CON}(N). \end{aligned}$$

Substituting $[1 - G(N+1)] = [1 - G(N)] - \delta_G(N)$ and $G(N+1) = G(N) + \delta_G(N)$, then grouping the resulting four terms by their structure, yields (8).

C Simulation Algorithm

Step 1. Construct $\mathcal{M}(N)$ per Definition 1, computing q_j and t_j for $j = 1, \dots, N$.

Step 2. Draw $n = 15,000$ consumers: $\theta_i \sim U[0, 1]$, $k_i = 1 + \text{Beta}(\alpha_G, \beta_G) \cdot 11$, independently.

Step 3. For each consumer i , compute $s_{ij} = \theta_i q_j - t_j$ for $j = 0, \dots, N$.

- If $k_i \geq N$ (sophisticated): $j^* = \min \arg \max_{j \in \{0, \dots, N\}} s_{ij}$.

- If $k_i < N$ (constrained): draw Bernoulli exit with probability $\gamma(N - k_i, k_i) = \min(1, \lambda(N - k_i)/(k_i + 1))$.
 - If exit: set $j^* = 0$.
 - Otherwise: draw uniform subset S of size $m = \min(\lfloor k_i \rfloor, N - 1)$ from $\{1, \dots, N\}$ without replacement. Draw $\varepsilon_{ij} \sim \mathcal{N}(0, \eta^2(N - k_i)^2/N^2)$ i.i.d. for each $j \in S$; set $\varepsilon_{i0} = 0$. Set $j^* = \min \arg \max_{j \in S \cup \{0\}} (s_{ij} + \varepsilon_{ij})$.

Step 4. Record: profit $t_{j^*} - C(q_{j^*})$ (if $j^* \geq 1$, else 0); participation ($j^* \geq 1$ and $q_{j^*} > 0$); efficiency ($j^* = j^{RAT}(\theta_i, N)$); surplus s_{ij^*} .

Step 5. Repeat for 10 seeds ($42 + r \cdot 1000 + N$, $r = 1, \dots, 10$) and report means.