

Token Is All You Price

Weijie Zhong*

This version: October 12, 2025

Abstract

We build a mechanism design framework where a platform designs GenAI models to screen users who obtain instrumental value from the generated conversation and privately differ in their preference for latency. We show that the revenue-optimal mechanism is simple: deploy a single aligned (user-optimal) model and use token cap as the only instrument to screen the user. The design decouples model training from pricing, is readily implemented with token metering, and mitigates misalignment pressures.

JEL: D47, D82, D83, L12, L86.

Keywords: screening, information design, large language models, token pricing, alignment.

1 INTRODUCTION

Motivation. Commercialization of GenAI has accelerated. During the exploration of the monetization strategy, usage-based (token) pricing has become a standard approach. Major providers publicly post per-million-token rates by model tier and meter the token usage.¹ On the other hand, there is much less consensus on how/whether to customize model design to cater to heterogeneous user needs. For example, OpenAI has been oscillating between offering a single model and offering a large menu of model variants, as is illustrated by [Table 1](#). This raises a fundamental design question: *How should a platform price and shape access to its models when users differ in their needs for the model?*

*Stanford GSB; email: weijie.zhong@stanford.edu

¹ e.g., OpenAI API pricing: <https://openai.com/api/pricing/>; Anthropic Claude pricing: <https://www.anthropic.com/pricing>; Google AI Studio/Gemini API pricing: <https://ai.google.dev/pricing>

Generation	Year	Models in the generation
1st	2022	GPT-3.5
2nd	2023	GPT-4
3rd	2024–2025	GPT-4o, GPT-4o mini, GPT-4.1, GPT-4.5, o1-preview, o1, o1-mini, o1-pro; o3, o3-mini, o3-pro; o4-mini
4th	2025	GPT-5, GPT-5 pro

Table 1: Historical versions of ChatGPT models offered.

Design challenge. There is huge heterogeneity among users’ preferences for models. For example, many assistant interactions are latency-sensitive (triage, hot-fixes, breaking news) while others are flexible (brainstorming, proof-reading). A user’s primary usage of the model determines his *urgency*, which is privately known by the user and varies among users. Classic mechanism design approach suggests offering a menu of customized products to screen the users (Mussa and Rosen, 1978; Maskin and Riley, 1984; Rochet and Choné, 1998). For GenAI, however, customization of models raises three challenges: Firstly, *the model* is a high-dimensional, complex object. Such complexity renders the mechanism design problem theoretically intractable. Secondly, scaling-law and compute-optimality results highlight the prohibitively high expense of pretraining a menu of customized models in practice (Kaplan et al., 2020; Hoffmann et al., 2022). Thirdly, aggressive profit-motivated customization can trigger a concern of *misalignment*: the models may behave differently from users preferred behaviors.²

This paper. We introduce a screening framework of GenAI model design and resolve these design tensions by solving for a simple revenue-maximizing mechanism. In our framework, a platform designs “(generative) models”, which generates the stochastic *conversation process*—summarized by the dynamic evolution of the user’s belief—subject to a per-time information cap that captures token throughput.³ A user interacts with the platform until he fully learns about an unknown state and collect a discounted payoff upon stopping. The user’s discount rate varies and is only known by the user.

² Broader AI-safety discussions warn about objective misspecification and reward hacking (Amodei et al., 2016; Ouyang et al., 2022; Bai et al., 2022).

³ Henceforth, we use “model”, “conversation process” and “belief process” in an interchangeable way. To avoid confusion, within the paper, we always use the noun “model” to refer to a GenAI model, not an economic model.

Our main result characterizes the optimal mechanism:

1. **Single aligned model.** The optimal mechanism uses a *uniform* (type-independent) model. Moreover, the model is *aligned* in the sense that it generates the user-optimal belief process if the user can run the model indefinitely.
2. **Exploratory conversation process.** The optimal GenAI model generates a simple “greedy” and “exploratory” belief process—probing the most promising approach (i.e. verifying the state that is closest to the current belief in the Bregman divergence.)—that occasionally yields decisive breakthroughs. The arrival rate of the breakthroughs is regulated by the token generation rate.
3. **Token price menu.** With a uniform model, the optimal mechanism screens users *solely* with a menu of token caps and prices. The token cap effectively imposes an additional deterministic stopping time on the conversation process—the model stops generating new information when the token limit is reached. The higher token cap, paired with higher price, targets higher user patience.

The result rationalizes prevalent token-based pricing and alignment-centered model training while explaining *why* a single general-purpose assistant paired with a tiered token menu is revenue-optimal in the presence of private information. The key intuition of the result hinges on the observation that the motive to screen users distorts the platform’s preference, i.e., the “virtual valuation” differs from user’s true valuation. However, under urgency heterogeneity, such distortion maintains the convexity of time preference up to a truncation of time. Therefore, the model design problem boils down to finding a uniformly optimal model for convex time preference, which is solved by the greedy exploration model. The screening problem boils down to truncating the time space, which is achieved by the token price menu.

Extensions. We study two extensions of the baseline framework. The first extension features additional heterogeneity on user valuation from learning. We show that the optimal mechanism remains a single aligned model and a token price menu when valuation is not too positively related to urgency. In the second extension, we endogenize the reasoning quality of the model by assuming the user obtains partial utility from imperfectly learning the state. We show that the opti-

mal menu features a menu of soft token caps: the reasoning quality deteriorates when the token cap is exhausted.

Relation to the literature (preview). We build on classic nonlinear pricing and screening (Mussa and Rosen, 1978; Maskin and Riley, 1984; Myerson, 1981; Armstrong, 1996; Rochet and Choné, 1998; Armstrong and Vickers, 2001; Laffont and Martimort, 2002; Bolton and Dewatripont, 2005); on Bayesian persuasion and dynamic information acquisition (Kamenica and Gentzkow, 2011; Che and Mierendorff, 2019; Che et al., 2023; Zhong, 2022; Sannikov and Zhong, 2024) and on rational inattention and information-theoretic constraints (Sims, 2003; Caplin and Dean, 2015; Matějka and McKay, 2015). The paper is closely related to the seminal paper of Bergemann et al. (2025) on LLM pricing, which we discuss in detail in [Section 5](#).

2 FRAMEWORK

An agent (*he*) would like to learn an unknown state $\theta \in \{1, \dots, n\}$ with prior $\mu_0 \in \Delta(\Theta)$. If the true state is learned in time $t \in \mathbb{R}_+$, the agent obtains utility e^{-rt} , where the discount rate $r \in [\underline{r}, \bar{r}]$ is the agent's privately known *urgency* parameter. The distribution of r has CDF G and PDF g , known to the principal. $0 < \underline{r} < \bar{r}$.

A principal (*she*) designs a GenAI model that generates a conversation with the user. The conversation induces a (càdlàg) posterior belief process $\langle \mu_t \rangle_{t \geq 0}$ about the state θ . $\langle \mu_t \rangle$ is a martingale due to the Bayes rule. The conversation ends when belief hits certainty:

$$\tau(\mu) = \inf\{t \geq 0 : \mu_t \in \{e_\theta\}_{\theta=1}^n\},$$

where $e_\theta = (\underbrace{0, \dots, 0}_{\theta-1}, 1, 0, \dots, 0)$ is the degenerate belief of state i . The conversation is subject to an *information throughput* (token-rate) constraint: for all $t, s \geq 0$,

$$\mathbb{E}[H(\mu_{t+s}) - H(\mu_t) \mid \mathcal{F}_t] \leq \chi s, \tag{Info.}$$

where $H : \Delta(\Theta) \rightarrow \mathbb{R}$ is a strictly convex (generalized entropy) function. The expected increasing in H measures the informational content generated, which is bounded by $\chi > 0$, the platform's token generation rate. Let \mathcal{M} denote the set of all càdlàg martingale processes that satisfy (Info.). For tractability, we assume that H is a $C^{(2)}$ smooth function and let ∇ and Hess denote the gradient and Hessian

operators, respectively.⁴ Let D be the Bregman divergence associated with H for $\mu, \mu' \in \Delta(\Theta)$:

$$D(\mu' | \mu) := H(\mu') - H(\mu) - \nabla H(\mu) \cdot (\mu' - \mu).$$

We impose the following regularity conditions on the Bregman divergence.

Assumption 1. $D(\cdot | \cdot)$ satisfies the following conditions

- $\sup_{\mu \in \Delta(\Theta)} \min_{\theta \in \Theta} D(e_\theta | \mu) < \infty$.
- There exists $\epsilon > 0$ s.t. $\forall \mu \in \Delta(\Theta)$ with $\mu(\theta) \leq \epsilon$, $D(e_\theta | \mu) > \min_{\theta' \neq \theta} D(e_{\theta'} | \mu)$.
- $\forall \mu \in \Delta(\Theta)^\circ, \forall \theta \neq \theta', (e_\theta - \mu) \nabla_\mu D(e_{\theta'} | \mu) \geq 0$.

Assumption 1 contains three parts. First, it states that the divergence from a belief to the “closest state” is bounded. This avoids the conversation process being “stuck” at belief at which no state can be learned. Second, it states that if a state is sufficiently unlikely under a belief, then the state is not the “closest state”. Third, it states that when μ moves closer to a state θ , the divergence to other states θ' does not decrease. The second and third conditions ensure that the divergence is consistent with the standard notion of “distance” and “direction” in the Euclidean space. **Assumption 1** is satisfied under canonical divergences like the KL divergence (derived from Shannon’s entropy) or the Mahalanobis divergence (derived from quadratic variation).

A (direct) mechanism specifies for each reported type r' (i) a model $\langle \mu_t^{r'} \rangle$ (with its probability space and filtration denoted by $(\Omega^r, \mathcal{P}^r, \mathcal{F}^r)$) and (ii) a transfer $P(r')$. Let $U(r' | r) = \mathbb{E}^{\mathcal{F}^r} [e^{-r\tau(\mu^{r'})}]$ denote the utility of true type r reporting r' . The principal chooses $\langle \mu_t^{r'} \rangle, P(\cdot)$ to maximize expected revenue subject to incentive compatibility (IC) and individual rationality (IR):

$$\max_{\langle \mu^{r'} \rangle \in \mathcal{M}, P(r)} \int_{\underline{r}}^{\bar{r}} P(r) g(r) dr \quad (\text{P})$$

$$\text{s.t. } U(r | r) - P(r) \geq \sup_{r'} U(r' | r) - P(r') \quad \forall r, \quad (\text{IC})$$

$$U(r | r) - P(r) \geq 0 \quad \forall r. \quad (\text{IR})$$

⁴ Throughout the paper, we normalize ∇H and $\text{Hess}H$ such that $\nabla H(\mu) \cdot \mu = H(\mu)$ and $\text{Hess}H(\mu)\mu = \mathbf{0}$. This normalization is obtained (without loss of generality) from extending H from $\Delta(\Theta)$ to $\mathbb{R}^{|\Theta|}$ via homogeneity of degree 1.

2.1 Discussion of modeling assumptions

Abstraction of the generated conversations. To study the design of GenAI models, we try to capture the most distinctive features of the transformer-based architecture (Vaswani et al., 2017): the model predicts the output tokens (keywords) of the conversation *sequentially, probabilistically* and at a *constant rate*. A generated conversation is then naturally model by a stochastic process with an information throughput constraint. The sequentiality of the generation also naturally calls for explicitly modeling the user’s preference for latency.

Given the complexity of the actual GenAI model implementations, we abstract away from a few other salient features. Most importantly, we do not model the user’s *input tokens* to the model, abstracting away from the *prompting* decision and corresponding pricing issues. We also abstract away from *sensing*—eliciting preference during the conversation—by assuming that communication about type only happens at $t = 0$.

Cost of production. We assume a zero cost of token production. Given that the computational cost of generation is almost linear in the output tokens (each output token corresponds to one full iteration of the neural network), one can introduce a per token cost $c \cdot t$ on the principal’s side. As it will be clear in our analysis, this will not change the structure of the optimal mechanism.

Other heterogeneity. User may be different in other dimensions. Another salient heterogeneity is user’s valuation for completing the task. We motivate the urgency heterogeneity using user’s primary usage of model in different tasks. However, different tasks are also of different importance of the user. We analyze an extension of our framework in [Section 4](#) to accommodate such heterogeneity by allowing the valuation to depend on urgency. We show that this will not change the structure of the optimal mechanism.

Endogenous reasoning quality. We assume that the model learns the state with certainty. This assumption is consistent with the prevalent practice of training the model against benchmarks — the models being shipped must reach a deterministic reasoning quality threshold. However, it is technically possible to train models using adaptive reasoning quality threshold, unlocking the possibility of screening users using the extra reasoning quality dimension. In [Section 4](#), we endogenize reasoning quality and illustrate how this changes the optimal model design.

2.2 Example

We illustrate our model in a simple example. Consider the setting where the agent’s task is summarized by learning a binary state $\theta = \{0, 1\}$ with prior belief $\frac{1}{2}$ denoting the probability of $\theta = 1$ being $\frac{1}{2}$. The agent’s privately known discount rate r is uniformly distributed on $[1, 2]$. The information throughput constraint is defined using the quadratic variation of the process: for all $t, s \geq 0$,

$$\mathbb{E}[(\mu_{t+s} - \mu_t)^2 \mid \mathcal{F}_t] \leq \frac{1}{8}s.$$

Per unit time, $\frac{1}{8}$ unit of token is generated; hence, the quadratic variation of $\langle \mu_t \rangle$ accumulates at rate $\leq \frac{1}{8}$. Then, the corresponding stopping time is $\tau = \inf\{t \mid \mu_t \in \{0, 1\}\}$ and the agent with type r obtains utility $\mathbb{E}[e^{-r\tau}]$ from the model $\langle \mu_t \rangle$.

Screening with constant-delay models. Assume for now that the platform only considers a parametric family of simple models that generate answers after a *constant* amount of time. Firstly, we determine such models that are feasible. Suppose the model learns the state deterministically at t , the posterior belief at t is either 0 or 1, each with $\frac{1}{2}$ probability. Therefore, the quadratic variation of μ_t is

$$\mathbb{E}[(\mu_t - 0.5)^2] = \frac{1}{2}(1 - 0.5)^2 + \frac{1}{2}(1 - 0.5)^2 \leq \frac{1}{8}t \iff t \geq 2,$$

where the inequality is implied by the information throughput constraint. In other words, a constant-delay model is feasible if it takes more than 2 unit of time to process.⁵ Next, we turn to the screening problem, where the principal sells a menu of such constant-delay models: $\{t(r), P(r)\}_{r \in [0, 1]}$, where $t(r) \geq 2$ is the delay and $P(r)$ is the transfer. The agent’s utility from reporting type r' is

$$U(r' \mid r) = e^{-rt(r')}.$$

Given the one dimensional family of models, the menu design problem reduces to the standard one-dimensional screening model of [Mussa and Rosen \(1978\)](#), which can be converted to the following unconstrained problem by a standard argument:

$$\sup_{t(r) \geq 2} \int_1^2 \underbrace{e^{-rt(r)} (1 - t(r)(r - 1))}_{\text{Virtual valuation of type } r} dr.$$

⁵ Note that we derived $t \geq 2$ as a necessary condition for feasibility. The sufficiency can be shown by explicitly constructing a belief process, which we omit in this example. See the “pure accumulation” strategy in [Chen and Zhong \(2025\)](#).

The optimal allocation is:

$$t^*(r) = \begin{cases} 2 & r \leq 1.5 \\ +\infty & r > 1.5. \end{cases}$$

That is to say, the optimal mechanism is a simple take-it-or-leave-it offer of the “efficient allocation” $t = 2$ at price $P = e^{-3}$. Types with discount rate greater than 1.5 declines the offer and types with discount rate less than 1.5 accepts the offer. The platform’s revenue is $\frac{1}{2}e^{-3} \approx 0.025$.

Screening with diffusion models. Now, we consider a different parametric model family that provides incremental information to the user during the conversation. Let $\langle \mu_t^\sigma \rangle$ be the Brownian motion defined by:

$$d\mu_t^\sigma = \sigma dB_t,$$

where $\langle B_t \rangle$ is a standard Brownian motion and $\mu_0^\sigma = \frac{1}{2}$. $\langle \mu_t^\sigma \rangle$ starts at the prior $\frac{1}{2}$ and diffuses with flow standard deviation σ . Since the flow variance is exactly the quadratic variation rate, $\langle \mu_t^\sigma \rangle$ satisfies the quadratic variation constraint if $\sigma^2 \leq \frac{1}{8}$. Qualitatively differently from the constant-delay models, the diffusion model $\langle \mu_t^\sigma \rangle$ learns the state at a stochastic time.

For any type r , the user’s utility from the diffusion model $\langle \mu_t^\sigma \rangle$ can be calculated analytically:⁶

$$U(\sigma | r) = \operatorname{sech}\left(\frac{\sqrt{r/2}}{\sigma}\right).$$

Therefore, the model design problem once again reduces to a standard one-dimensional screening problem, which can be solved by analyzing the unconstrained problem:

$$\sup_{\sigma(r) \leq \frac{1}{2\sqrt{2}}} \int_1^2 \operatorname{sech}\left(\frac{\sqrt{r/2}}{\sigma(r)}\right) \left(1 - \frac{r-1}{2\sqrt{2}r\sigma(r)} \tanh\left(\frac{\sqrt{r/2}}{\sigma(r)}\right)\right) dr.$$

The optimal allocation is $\sigma^*(r) \equiv \frac{1}{2\sqrt{2}}$, i.e., the optimal mechanism is to sell the “efficient allocation” $\sigma = \frac{1}{2\sqrt{2}}$ to all types, which determines the revenue being $\operatorname{sech}(2\sqrt{2}) \approx 0.12$.

⁶ $U(\sigma | r)$ is derived from solving the HJB equation $rV(\mu) = \frac{1}{2}\sigma^2 V''(\mu)$ with initial conditions $V(0) = V(1) = 1$ and evaluating the solution at $V(0.5)$. See Bolton and Harris (1999).

In both examples, the optimal mechanism is to sell a model at a take-it-or-leave-it price and the diffusion model clearly outperforms the constant-delay model in terms of revenue. This seems to suggest that model design is the central problem for revenue maximization while the pricing strategy tends to be simple. However, optimizing the model design is an unconventional mechanism design problem with a very rich space of screening instruments, which features the theoretical and practical challenges we discussed earlier in the introduction. In the next section, we fully solve the mechanism design problem and show that the optimal design is exactly the opposite of what the examples suggested: it features a simple model but a more sophisticated pricing strategy.

3 ANALYSIS

In this section, we solve (P) through three steps. We first reduce the constrained optimization problem to a relaxed pointwise optimization problem by considering only local IC's. Secondly, we show that the solution to the pointwise optimization problem is a type-independent belief process (termed a greedy exploration model) truncated at a type-dependent time (token cap). Then, we verify that the mechanism that employs the greedy exploration model and screens the user using token caps is optimal.

Reduction to pointwise problem. Fix a menu $(\langle \mu_t^r \rangle, P(r))$. Following the standard envelope theorem argument, (IC) implies that

$$\begin{aligned} \frac{dU(r | r) - P(r)}{dr} &= \left. \frac{\partial U(r' | r)}{\partial r} \right|_{r'=r} = \mathbb{E}^{\mathcal{P}^r} [-e^{-r\tau(\mu^r)} \tau(\mu^r)]. \\ \implies P(r) &= \mathbb{E}^{\mathcal{P}^r} [e^{-r\tau(\mu^r)}] - \int_{\bar{r}}^r \mathbb{E}^{\mathcal{P}^z} [-e^{-z\tau(\mu^z)} \tau(\mu^z)] dz - U(\bar{r} | \bar{r}) + P(\bar{r}). \end{aligned}$$

By setting $U(\bar{r} | \bar{r}) - P(\bar{r}) = 0$, we obtain a relaxed revenue object;

$$\sup_{\langle \mu_t^r \rangle \in \mathcal{M}} \int_{\underline{r}}^{\bar{r}} \mathbb{E}^{\mathcal{P}^r} \left[e^{-r\tau(\mu^r)} \left(1 - \tau(\mu^r) \frac{G(r)}{g(r)} \right) \right] g(r) dr. \quad (1)$$

Note that (1) is separable in r ; hence, it can be solved point-wise. Next, we consider the pointwise optimization problem that maximize the virtual value of the model:

$$\sup_{\langle \mu_t^r \rangle \in \mathcal{M}} \mathbb{E}^{\mathcal{P}^r} \left[e^{-r\tau(\mu^r)} \left(1 - \tau(\mu^r) \frac{G(r)}{g(r)} \right) \right]. \quad (2)$$

We call the integrand $e^{-rt} \left(1 - t \frac{G(r)}{g(r)}\right)$ the “virtual time preference” of the principal. The virtual time preference is the user’s time preference adjusted by the time-dependent information rent.

Optimal model design. We solve (2) by verifying a conjectured solution. We consider the following *greedy exploration model*.

Definition 1. The *greedy exploration model* $\langle \mu_t^* \rangle$ is defined recursively. The initial parameters are: $k = 1$, $\widehat{\mu}^1 = \mu_0$, $\widehat{t}^1 = 0$ and $\Theta^1 = \operatorname{argmin}_{\theta \in \Theta} D(e_\theta | \widehat{\mu}^1)$.

- While $\Theta^k \subsetneq \Theta$, for $t \geq \widehat{t}^k$, define $(\beta_t(\theta) \geq 0)_{\theta \in \Theta^k}$ and $\widehat{\mu}_t$ via the following functional equations:

$$D(e_\theta | \widehat{\mu}_t) = D(e_{\theta'} | \widehat{\mu}_t), \quad \forall \theta, \theta' \in \Theta^k; \quad (3)$$

$$\sum_{\theta \in \Theta^k} \beta_t(\theta) D(e_\theta | \widehat{\mu}_t) = \chi; \quad (4)$$

$$\frac{d\widehat{\mu}_t}{dt} = - \sum_{\theta \in \Theta^k} \beta_t(\theta) (e_\theta - \widehat{\mu}_t); \quad (5)$$

$$\widehat{\mu}_t \Big|_{t=\widehat{t}^k} = \widehat{\mu}^k. \quad (6)$$

Let \widehat{t}^{k+1} be the earliest time when $\Theta^k \subsetneq \operatorname{argmin}_{\theta \in \Theta} D(e_\theta | \widehat{\mu}_t)$. Let $\widehat{\mu}^{k+1} = \widehat{\mu}_{\widehat{t}^{k+1}}$ and $\Theta^{k+1} = \operatorname{argmin}_{\theta \in \Theta} D(e_\theta | \widehat{\mu}^{k+1})$.

Repeat the iteration with $k = k + 1$.

- If $\Theta^k = \Theta$, for $t \geq \widehat{t}^k$, define $\widehat{\mu}_t \equiv \widehat{\mu}^k$ and $\beta_t(\theta) \equiv \widehat{\mu}^k(\theta)$. Let $K = k$ and $\widehat{t}^{K+1} = \infty$. End the iteration.

For $t \in [\widehat{t}^k, \widehat{t}^{k+1})$, $\langle \widehat{\mu}_t^* \rangle$ is the following compensated Poisson process:

$$d\mu_t^* = \sum_{\theta \in \Theta^k} \left(dQ_t^\theta(\beta_t(\theta)) - \beta_t(\theta) dt \right) (e_\theta - \mu_t^*),$$

where $Q_t^\theta(x)$ are independent Poisson counters with arrival rate x .

While **Definition 1** seems complicated, it describes a very simple model. At every moment in time, given the current belief μ_t , the model allocates its ‘token-rate budget’ across currently closest states (in Bregman divergence), i.e., states in set Θ^k in **Definition 1**. Either a decisive jump to a state occurs (a breakthrough) at Poisson

rate β_t , or, absent a jump, the posterior drifts away so that new states eventually enter the ‘closest set’.

The Poisson rates β_t are pinned down by two conditions. Firstly, the information throughput is exhausted (Equation (4)). Secondly, the continuing belief process maintains the same divergence to each of the states in Θ_k (Equation (3)). The compensating drift is pinned down by Bayes rule (Equation (5)). For sufficiently large $t \geq t^K$, all states enter in the consideration set and the belief process becomes stationary.

The name “greedy exploration model” is intuitive: the model maximizes the arrival rate of an instantaneous revealing signal about some state. Therefore, the greedy exploration model is optimal for a myopic user. Since the greedy model is myopic, it ignores the negative consequence that absent the Poisson jump, the subsequent task of learning the state becomes more difficult and takes longer to process.

A key auxiliary result Proposition 1 shows that the greedy exploration model is optimal under any payoff function that is convex in the stopping time.

Proposition 1. *The greedy exploration model $\langle \mu_t^* \rangle$ is (uniquely) well-defined and for all positive, decreasing, convex and continuous function $\rho(t)$,*

$$\langle \mu_t^* \rangle \in \arg \max_{\langle \mu_t \rangle \in \mathcal{M}} \mathbb{E}^{\mathcal{P}} [\rho(\tau(\mu))]. \quad (7)$$

Compare (2) and (7), the “virtual time preference” in (2) for type r is $\rho(t) = e^{-rt}(1 - tG(r)/g(r))$. Let $T(r) = \frac{g(r)}{G(r)}$. $\rho(t)$ is strictly convex and positive for $t < T(r)$ and strictly negative for $t > T(r)$. Since ρ is not globally convex, Proposition 1 does not directly apply. However, $\rho(t)$ can be truncated at 0 to be converted to a convex function. Formally, let $\rho(t)^+ := \max\{\rho(t), 0\} = \rho(\max\{t, T(r)\})$. Then, ρ^+ is decreasing and weakly convex; hence, Proposition 1 implies that $\langle \mu_t^* \rangle$ maximizes $\mathbb{E}[\rho(\tau(\mu))^+]$. Note that $\mathbb{E}[\rho(\tau(\mu^*))^+]$ is achieved by the $T(r)$ truncation of $\langle \mu_t^* \rangle$: define

$$\mu_t^{*T} := \mu_{\min\{t, T\}}^*$$

Then, $\langle \mu_t^{*T(r)} \rangle$ solves (2). To complete our analysis, we show that the value of the relaxed problem (2) is achieved by an IC and IR mechanism. The proof of Proposition 1 is deferred to the end of the section.

Optimal menu and token pricing. Consider the following menu

Definition 2. The *token price menu* $(\langle \mu_t^* \rangle, \{(\chi T(r), P(r))\}_{r \in [L, \bar{r}]})$ consists of

- A single greedy exploration model $\langle \mu_t^* \rangle$, with corresponding stopping time $\tau(\mu^*) \sim f^*$.
- A menu $\{(\chi T(r), P(r))\}_{r \in [L, \bar{r}]}$, where

$$\begin{aligned} \text{Token cap: } \chi T(r) &= \frac{\chi g(r)}{G(r)}; \\ \text{Price: } P(r) &= \int_0^{T(r)} e^{-rt} f^*(t) dt - \int_r^{\bar{r}} \int_0^{T(z)} e^{-zt} f^*(t) dt dz. \end{aligned}$$

The marginal price of the token cap (per extra token) can be calculated by:

$$\frac{P^{*'}(r)}{T'(r)} = e^{-rT(r)} f(T(r))$$

and is decreasing in $\chi T(r)$, the token consumption.

Theorem 1. The token price menu $(\langle \mu_t^* \rangle, \{(\chi T(r), P(r))\}_{r \in [L, \bar{r}]})$ is an optimal mechanism.

Proof. By [Proposition 1](#), the token price menu achieves the value of the relaxed pointwise optimization problem (1). We verify that it satisfies (IC) and (IR) globally. Since the menu is one-dimensional, $P(r)$ is derived from the local (IC) and (IR) is satisfied for $r = \bar{r}$, it is sufficient to check the supermodularity condition for $U(r' | r)$:

$$\begin{aligned} & \frac{\partial^2}{\partial r \partial r'} \int_0^{T(r')} e^{-rt} f^*(t) dt \\ &= -T'(r') T(r') e^{-rT(r')} f^*(t) \\ &\geq 0. \end{aligned}$$

Q.E.D.

[Theorem 1](#) states that the optimal mechanism consists of the greedy exploration model and a token price menu. The first and perhaps most striking feature of our prediction is that the optimal menu of models offered is a single model. Moreover, a direct implication of [Proposition 1](#) is that the optimal model is “aligned”

in the sense that it maximizes the user surplus. Therefore, the model design problem is effectively decoupled from the mechanism design problem. The existing paradigm that trains the GenAI model based on “matching human behavior” turned out to be the optimal approach for revenue maximization.

Secondly, the optimal model features “greedy exploration”, which is preferred under any time-convex payoff function. An important insight from our analysis is that the principal’s virtual time preference has a strictly convex positive part:

$$e^{-r} \left(1 - t \frac{G(r)}{g(r)} \right)$$

and the negative part can be truncated to 0 as the principal can stop the model at any time. That is to say, under urgency heterogeneity, the incentive adjusted virtual time preference is aligned with the agent’s exponential discounting time preference. This convexity renders greedy exploration a uniformly optimal model for both the principal and the agent, as the greedy Poisson process maximizes the dispersion of stopping time by generating early stops as well as long delays.⁷

Thirdly, the optimal menu is a token price menu. For every price $P(r)$, the users obtains a token cap of $\chi T(r)$. The model provides meaningful outputs until the token cap is exhausted. Such price scheme is consistent with the practice in industry: platforms typically provides different tiers of token caps. For more demanding commercial users, the platform may further customize the token usage contract with a more detailed price schedule.

3.1 Example revisited

It is easy to verify that the divergence from 0.5 to the two states are the same:

$$D(0 | 0.5) = D(1 | 0.5) = \frac{1}{4}.$$

Therefore, the *greedy exploration model* generates a simple stationary Poisson process, where belief either stays at 0.5 or jumps to 0 or 1, each with rate $\frac{\chi}{1/4} = \frac{1}{2}$. The corresponding stopping time is exponential:

$$f^*(t) = \frac{1}{2} e^{-\frac{1}{2}t}.$$

⁷ The connection between time preference and the optimality of greedy exploration has been discussed in [Chen and Zhong \(2025\)](#) in a simpler binary state setting. Our [Proposition 1](#) can be viewed as an extension of [Chen and Zhong \(2025\)](#) to a general state space.

For each type r , the token cap $\chi T(t) = \chi \frac{g(r)}{G(r)} = \frac{1}{8} \frac{1}{r-1}$. The price can be calculated correspondingly:

$$\begin{aligned} P(r) &= \int_0^{\frac{1}{r-1}} \frac{1}{2} e^{-(r+1/2)t} dt - \int_r^2 \int_0^{\frac{1}{z-1}} \frac{1}{2} e^{-(z+1/2)t} t dt dz \\ &= \frac{1}{15} \left(3 + 2e^{-5/2} - 5e^{-\frac{r+1/2}{r-1}} \right). \end{aligned}$$

The optimal revenue of the principal:

$$\begin{aligned} \pi^* &= \int_1^2 \int_0^{\frac{1}{r-1}} \frac{1}{2} e^{-(r+1/2)t} (1 - t(r-1)) dt dr \\ &= 0.2(1 - e^{-2.5}) + 0.5e^{-1} \int_{1.5}^{\infty} e^{-z}/z dz \\ &\approx 0.2. \end{aligned}$$

The optimal mechanism generates 8 times the revenue of the constant-delay model family and 2 times the revenue of the diffusion model family.

3.2 Proof of Proposition 1

Proof. Step I. We begin by verifying that $\langle \mu_t^* \rangle$ is well-defined according to Definition 1. Inductively we assume that the process has been defined for round $k-1$, with $\widehat{\mu}^k \in \Delta(\Theta)^\circ$. Consider the construction in round k . By the definition of Θ^k , $\forall \theta, \theta' \in \Theta^k$, $D(e_\theta | \widehat{\mu}^k) = D(e_{\theta'} | \widehat{\mu}^k)$; hence, (3) is equivalent to

$$\begin{aligned} \frac{d}{dt} D(e_\theta | \widehat{\mu}_t) &= \frac{d}{dt} D(e_{\theta'} | \widehat{\mu}_t) \\ \iff (e_\theta - \widehat{\mu}_t)^\top \text{Hess}H(\widehat{\mu}_t) \frac{d\widehat{\mu}_t}{dt} &= (e_{\theta'} - \widehat{\mu}_t)^\top \text{Hess}H(\widehat{\mu}_t) \frac{d\widehat{\mu}_t}{dt} \\ \iff (e_\theta - e_{\theta'})^\top \text{Hess}H(\widehat{\mu}_t) \left(\sum_{\theta'' \in \Theta^k} \beta_t(\theta'') (e_{\theta''} - \widehat{\mu}_t) \right) &= 0. \end{aligned}$$

Combining the equations for all pairs of $\theta, \theta' \in \Theta^k$, we get

$$\begin{aligned} \left[e_\theta - e_{\theta'} \right]_{\theta, \theta' \in \Theta^k}^\top \cdot \text{Hess}H(\widehat{\mu}_t) \cdot \left[e_{\theta''} - \widehat{\mu}_t \right]_{\theta'' \in \Theta^k} \cdot \beta_t &= 0 \\ \iff \left[e_\theta - e_{\theta'} \right]_{\theta, \theta' \in \Theta^k}^\top \cdot \text{Hess}H(\widehat{\mu}_t) \cdot \left[e_{\theta''} \right]_{\theta'' \in \Theta^k} \cdot \beta_t &= 0, \\ \iff \left[\text{Hess}H(\widehat{\mu}_t)_{\theta, \theta'} \right]_{\theta, \theta' \in \Theta^k} \cdot \beta_t &= \alpha \mathbf{1}, \quad \text{for some } \alpha \neq 0. \end{aligned} \tag{8}$$

where the first equality is implied by $\text{Hess}H(\widehat{\mu}_t) \cdot \widehat{\mu}_t = 0$; the second equality is from the fact $e_\theta \text{Hess}H(\widehat{\mu}_t) e_{\theta'} = \text{Hess}H(\widehat{\mu}_t)_{\theta, \theta'}$. Let $\Sigma(\mu) = [\text{Hess}H(\mu)_{\theta, \theta'}]_{\theta, \theta' \in \Theta^k}$ be submatrix of $\text{Hess}H$ restricted to the indices in Θ , (8) can be written as

$$\Sigma(\widehat{\mu}_t) \cdot \beta_t = \alpha \mathbf{1}.$$

Note that $\text{Hess}H(\mu)$ has a unique eigenvector μ corresponding to eigenvalue 0. Therefore, $\forall \mu \in \Delta(\Theta)$, $(\mu(\theta) \mathbf{1}_{\theta \in \Theta^k})_{\theta \in \Theta}^\top \text{Hess}H(\widehat{\mu}_t) \neq 0$ since $\widehat{\mu}_t$ is interior, and $|\Theta^k| < n$. Therefore $\Sigma(\widehat{\mu}_t)$ is invertible and $\beta = \alpha \Sigma(\widehat{\mu}_t)^{-1} \cdot \mathbf{1}$.

Part 3 of **Assumption 1** is equivalent to $(e_\theta - \mu)^\top \text{Hess}H(\mu)(e_{\theta'} - \mu) \leq 0$, implying that $\text{Hess}H(\mu)$ is a Stieltjes matrix. Then, $\Sigma(\widehat{\mu}_t)$ is also a positive definite Stieltjes matrix. As a result, the vector $\Sigma(\widehat{\mu}_t)^{-1} \mathbf{1}$ is strictly positive. (4) then uniquely pins down the positive scalar α and the vector β :

$$\beta_t = \frac{\chi \Sigma(\widehat{\mu}_t)^{-1} \cdot \mathbf{1}}{[D(e_\theta | \widehat{\mu}_t)]_{\theta \in \Theta^k}^\top \cdot \Sigma(\widehat{\mu}_t)^{-1} \cdot \mathbf{1}}.$$

Since the eigenvalues of $\Sigma(\widehat{\mu}_t)$ are bounded away from zero, the inversion operation is Lipchitz continuous. Then, we obtained β_t as a Lipchitz continuous function of $\widehat{\mu}_t$. This, together with the differential equation (5) uniquely pins down a $C^{(1)}$ smooth solution $\widehat{\mu}_t$, with initial condition (6) by the Picard-Lindelof theorem.

Next, we show that \widehat{t}^{k+1} exists. Note that $\max_\theta \beta_t(\theta)$ is bounded below by $\delta = \frac{\chi}{\sup_\mu \min D(e_\theta | \mu)} > 0$ by part 1 of **Assumption 1**. That is to say, there exists finite time t' at which $\widehat{\mu}_{t'}(\theta)$ will fall below ϵ (define in part 2 of **Assumption 1**) for some $\theta \in \Theta^k$. This suggests that θ is no longer the closest state to $\widehat{\mu}_{t'}$ under the Bregman divergence. **Equation (3)** implies that all states in Θ^k are equally far from $\widehat{\mu}_t$ for all $t \geq t^k$; hence, there exists an earliest time t^{k+1} that a state outside of Θ^k enters $\arg \min D(e_\theta | \widehat{\mu}_t)$ before $\widehat{\mu}_t(\theta)$ fall below ϵ for some θ (by the continuity of the path $\widehat{\mu}_t$ and the divergence D). This also proves that $\widehat{\mu}^{k+1} \in \Delta(\Theta)^\circ$.

Finally, since the number of states is finite, the iteration ends in finite time, with $\widehat{\mu}^{K-1} \in \Delta(\Theta)^\circ$. Therefore, we obtain an interior path $\widehat{\mu}_t$ defined on \mathbb{R}_+ and strictly positive $\beta_t(\theta)$ defined on $[\widehat{t}^k, \infty)$, where k is the smallest index that $\theta \in \Theta^k$. $\widehat{\mu}_t$ is continuous and piecewise $C^{(1)}$ on each interval $(\widehat{t}^k, \widehat{t}^{k+1})$.

Step II. Next, we simplify (7) by reducing it to a linear program. Any feasible conversation $\langle \mu_t \rangle$ that satisfies (**Info.**) induces a distribution of stopping times defined as follows. Let F^i denote the CDF of $\tau(\mu)$ conditional on state $\theta = i$ (with

density f^i). Let \mathbf{F} denote $(F^i)_{i=1}^n$, Let $F = \sum_i F^i$ (with density f). (Sannikov and Zhong, 2024, Theorem 1) provides a complete characterization of such \mathbf{F} : $\langle \mu_t \rangle$ satisfies (Info.) if and only if for all $t \geq 0$,

$$\sum_i F^i(t)H(e_i) + H\left(\sum_i (F^i(\infty) - F^i(t))e_i\right) - H(\mu_0) \leq \chi \int_0^t (1 - F(s)) ds. \quad (9)$$

Therefore, (7) is equivalent to

$$\begin{aligned} \sup_{\mathbf{F}} \int_0^\infty \rho(t)f(t) dt \\ \text{s.t. (9).} \end{aligned} \quad (10)$$

Step III. For now, we assume in addition that ρ is strictly convex and $\mathbf{C}^{(2)}$. We verify that the stopping time distribution \mathbf{f}^* induced by model $\langle \mu_t^* \rangle$ solves (10). By (Sannikov and Zhong, 2024, Theorem 2), it is sufficient to find $a \in \mathbb{R}^{|\Theta|}$ and positive measure $\lambda \in L^1$ to establish the first order condition:

$$l_{\mathbf{f}^*, \lambda}(\theta, t) = \rho(t) + \chi \int_{s \leq t} \Lambda(t) ds - \int_{s \in (0, t)} \nabla H(\widehat{\mu}_t) d\lambda(s) \cdot e_\theta - \Lambda(t)H(e_\theta) \leq a \cdot e_\theta, \quad (11)$$

where $\Lambda(t) = \int_t^\infty \lambda(ds)$, with equality holding on the support of \mathbf{f}^* and complementary slackness conditions satisfied. Note that Sannikov and Zhong (2024) defines l on the space of all stopping beliefs. Here, since the stopping beliefs are e_θ 's only, we directly define l as a function of θ . We establish the FOC by explicitly constructing a and λ . Let

$$\zeta(t) = \min_{\theta} D(e_\theta | \widehat{\mu}_t).$$

Recall that $\zeta(t)$ is achieved by every $\theta \in \Theta^k$ for $t \geq \widehat{t}^k$. By construction, $\zeta(t)$ is bounded away from 0. By the continuity of $\widehat{\mu}_t$, $\zeta(t)$ is Lipschitz continuous. The following ODE

$$\frac{d\Lambda(t)}{dt} = \frac{\rho'(t) + \chi\Lambda(t)}{\zeta(t)},$$

with initial condition $\Lambda(+\infty) = 0$ has an explicit solution $\Lambda(t)$:

$$\Lambda(t) = \int_t^\infty \frac{e^{\int_s^t \frac{\chi}{\zeta(z)} dz} (-\rho'(s))}{\zeta(s)} ds.$$

We verify that Λ is strictly decreasing, i.e., λ is a positive function. The ODE implies

$$\rho''(t) + \lambda(t)(\zeta'(t) - \chi) = \zeta(t)\lambda'(t).$$

Therefore, when $\lambda(t) = 0$, $\lambda'(t) > 0$ due to the strict convexity of ρ . Hence, λ never crosses 0 and $\lambda(t) = -\Lambda'(t) > 0$ when $t \rightarrow \infty$. Therefore, $\lambda > 0$ as desired. Given the constructed Λ , $\forall \theta$ and t ,

$$\begin{aligned} \frac{dl_{f^*,\lambda}(\theta, t)}{dt} &= \rho'(t) + \chi\Lambda(t) + \lambda(t)D(e_\theta | \widehat{\mu}_t) \quad \begin{cases} = 0 & \text{if } \theta \in \Theta^k \text{ and } t \geq \widehat{t}^k \\ > 0 & \text{otherwise} \end{cases} \\ \Rightarrow l_{f^*,\lambda}(\theta, t) &\begin{cases} = l_{f^*,\lambda}(\theta, \widehat{t}^k) & \text{if } \theta \in \Theta^k \text{ and } t \geq \widehat{t}^k \\ < l_{f^*,\lambda}(\theta, \widehat{t}^k) & \text{otherwise.} \end{cases} \end{aligned}$$

Therefore, let $a_\theta = l_{f^*,\lambda}(\theta, \widehat{t}^k)$, the FOC (11) is satisfied. Note that the constraint (Info.) is binding all the time by construction (see (4)). Hence, the complementary slackness condition trivially holds. Therefore, we conclude that

$$\langle \mu_t^* \rangle \in \arg \max_{\langle \mu_t \rangle} \mathbb{E}^{\mathcal{P}}[\rho(\tau(\mu))].$$

Finally, for a general ρ , $\forall \eta > 0$, a standard mollification argument suggests that there exists $\rho_\eta \in [\rho, \rho + \eta]$ and ρ_η is strictly convex and $C^{(2)}$ smooth (see, e.g., (Aza-gra and Stolyarov, 2023, Theorem 5)). The argument above implies that $\langle \mu_t^* \rangle$ maximizes $\mathbb{E}[\rho_\eta(\tau(\mu))]$. Then,

$$\begin{aligned} \sup_{\langle \mu \rangle \in \mathcal{M}} \mathbb{E}[\rho(\tau(\mu))] &\leq \mathbb{E}[\rho_\eta(\tau(\mu^*))] \leq \mathbb{E}[\rho_\eta(\tau(\mu^*))] + \eta \\ \xrightarrow{\eta \rightarrow 0} \sup_{\langle \mu \rangle \in \mathcal{M}} \mathbb{E}[\rho(\tau(\mu))] &= \mathbb{E}[\rho_\eta(\tau(\mu^*))]. \end{aligned}$$

Q.E.D.

4 EXTENSIONS AND ROBUSTNESS

Heterogeneous valuations. Consider a setting where users also differ in their valuation for the output. Specifically, we assume that user's utility from learning in period t is $q(r)e^{-rt}$, where r is still privately known to the user only. The scalar

$q(r)$ is $C^{(1)}$. Following the same envelope theorem argument, the “virtual time preference” of the principal for type r is

$$q(r)e^{-rt} \left(1 - \left(t - \frac{q'(r)}{q(r)} \right) \frac{G(r)}{g(r)} \right).$$

It is easy to verify that the virtual time preference is strictly positive and convex for $t < T(r) = \frac{g(r)}{G(r)} + \frac{q'(r)}{q(r)}$ and strictly negative for $t > T(r)$. Therefore, as long as the crossing point $T(r)$ is decreasing in r , the solution to the pointwise optimization problem is the same as the baseline framework, except that $T(r)$ might be negative for sufficiently large r 's, meaning that those types are excluded from the market. Then, the supermodularity condition can be verified:

$$\frac{\partial^2}{\partial r \partial r'} U(r | r') = T'(r')(q'(r) - T(r')q(r))e^{-rT(r')} \sim T(r') - \frac{q'(r)}{q(r)}. \quad (12)$$

Evidently, when $q' < 0$, [Equation \(12\)](#) is positive. More generally, when $\max_r \frac{q'(r)}{q(r)} < T(\bar{r})$, i.e., q is not too steeply increasing, [Equation \(12\)](#) is positive.

To sum up, value heterogeneity does not change our prediction. The optimal mechanism is still a user-optimal model and a token price menu.

Endogenous reasoning quality. Consider the case where the principal can also choose “reasoning quality” over the conversation. Assume now that the user obtains utility $V(\mu)$ if the model stops at belief μ . $V(\mu)$ captures the user’s utility from reasoning quality of the model. The principal now controls the endogenous stopping time as well. In this case, the virtual preference of the principal is

$$V(\mu)e^{-rt} \left(1 - t \frac{G(r)}{g(r)} \right).$$

The principal’s pointwise optimization problem solves

$$\max_{\langle \mu_t, \tau \rangle} \mathbb{E}^{\mathcal{P}} \left[V(\mu_\tau) e^{-r\tau} \left(1 - \tau \frac{G(r)}{g(r)} \right) \right]. \quad (13)$$

(13) is generally a much more difficult problem than (10). To obtain intuition, we consider a simple parametric setting: $V(\mu) = |\mu - 0.5|$ and $H(\mu) = |\mu - 0.5|^\alpha$. Due to the result of [Sannikov and Zhong \(2024\)](#), the problem has an analytical solution: the reasoning quality $\kappa = |\mu - 0.5|$ is given by a decreasing function

$$\kappa(t) = \left(\frac{\alpha \chi}{(\alpha - 1)(\alpha + 1)} (T(r) - t) e^{-\alpha r (T(r) - t)} {}_1F_1(\alpha + 1, \alpha + 2, \alpha r (T(r) - t)) \right)^{\frac{1}{\alpha}},$$

where ${}_1F_1$ is the hypergeometric function and $T(t) = \frac{g(r)}{G(r)}$. In the limit $T \rightarrow \infty$, $\kappa \rightarrow \left(\frac{\chi}{(\alpha-1)r}\right)^{\frac{1}{\alpha}}$, a constant reasoning quality. In **Figure 1**, we plot the stopping boundaries $0.5 \pm \kappa(t)$ for a set of r using the same parameter as our leading example: $\chi = 1/8$ and $\alpha = 2$.

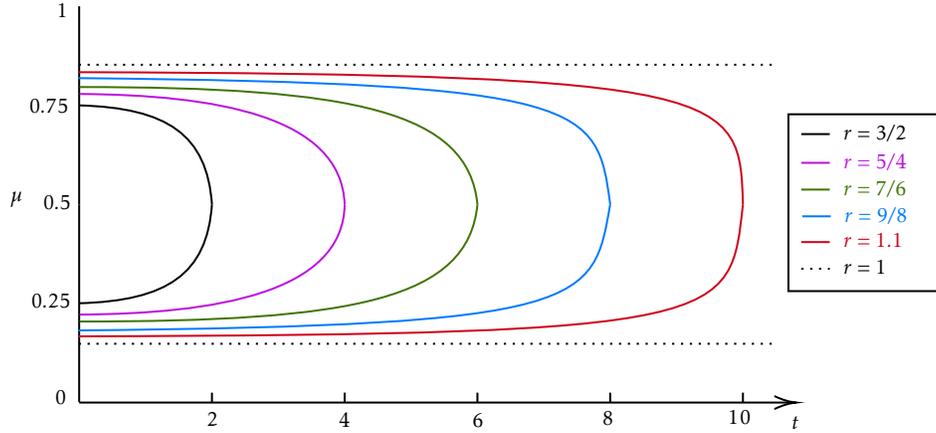


Figure 1: Token consumption based reasoning quality.

The solution exhibits two new features. Firstly, there is a quality-delay trade-off: more urgent types obtain overall lower reasoning quality but higher stopping rate. Secondly, for all types except for \underline{r} , the reasoning quality gradually decays to zero when t gets to $T(r)$. The most patient type continues to obtain the stationary exploratory mode that is user-optimal. This restores a role for limited type-contingent model design. However, the optimal menu of models can still be implemented by a relatively simple design, e.g., a single aligned pretrained model plus a series of post trained deteriorating factors (fine tuning) that limits the models' reasoning quality based on token consumption, leveraging the relatively cheap fine tuning procedure.

5 RELATED LITERATURE

Our paper sits at the intersection of *screening and nonlinear pricing*, *dynamic information acquisition and persuasion*, *rational inattention and information-theoretic constraints*, and the *economics and computer systems of large language models (LLMs)*. We connect these literatures by studying a monopolist who designs and prices GenAI models under an information-throughput (token) constraint and show that a single aligned model and a menu of token caps maximizes revenue.

Screening, nonlinear pricing, and mechanism design. We build on the classic second-degree price discrimination literature in which a monopolist designs menus to screen heterogeneous buyers (Mussa and Rosen, 1978; Maskin and Riley, 1984; Myerson, 1981). Our screening instrument is unconventional: rather than product quality or quantity, we screen via the *distribution of stopping times* generated by a conversation constrained by information flow. Multidimensional screening analyses such as Armstrong (1996) and Rochet and Choné (1998), and competitive variants such as Armstrong and Vickers (2001), illuminate the difficulty of screening when products are high-dimensional.

Mechanism design for the domain of GenAI models is an emerging field. A closely related pioneering paper Bergemann et al. (2025) ask the same central question as our paper and study the optimal screening mechanisms for selling GenAI models. Bergemann et al. (2025) model the screening instruments in a more reduced-form way—a GenAI model is a scalar output function that takes token usage and fine-tuning parameter as input—and focus on screening multi-dimensional heterogeneity. Our framework differs by modeling the generative models explicitly using the generated conversation process and focus on model design under (one-dimensional) heterogeneous preference over latency. The methodologies of the two papers exactly complement each other: Bergemann et al. (2025) can be viewed as a framework of screening via *model post-training* and our paper can be viewed as a framework of screening via *model pre-training*.

Dynamic information acquisition and persuasion. We treat a generated conversation as a designed stochastic process over beliefs that ends at a stopping time. Methodologically, we leverage results that characterize implementable stopping distributions under controlled exploration, most directly Sannikov and Zhong (2024). The connection between convex time preference and greedy exploration has been established in Chen and Zhong (2025) in a simpler binary state setting. Related strands include speed–accuracy tradeoffs and optimal stopping (Fudenberg et al., 2018; Morris and Strack, 2019), dynamic attention allocation across sources (Che and Mierendorff, 2019), and persuasion with limited patience where a sender must keep the receiver engaged (Che et al., 2023). We embed a common optimal learning problem from this literature into a mechanism design setting.

Our feasibility constraint—bounded entropy reduction per unit time—echoes

the information-channel constraints pioneered by Sims (2003) and subsequent RI work that micro-founds choice under attention limits (Caplin and Dean, 2015; Matějka and McKay, 2015; Bloedel and Zhong, 2024).

Alignment and incentive distortions. Empirical and theoretical work shows that optimizing for engagement or profit may misalign systems. In a preregistered audit of Twitter/X, Milli et al. (2025) find that engagement ranking amplifies divisive content and departs from users’ tweet-level stated preferences; a large randomized experiment that moved Facebook and Instagram users to chronological feeds reduced activity without short-run attitude change, indicating engagement is a poor proxy for welfare (Guess et al., 2023). At the organizational level, capability–safety trade-offs (“alignment tax”) imply that post-training for human preference/safety can lower measured performance, creating incentives to relax alignment under market pressure (Lin et al., 2024). By decoupling *alignment* (a single model trained for user welfare) from *monetization* (token caps and prices), our mechanism reduces pressure to personalize the model itself, mitigating misalignment risks while allowing efficient screening via usage.

6 CONCLUSION AND DISCUSSION

We characterize the revenue-optimal way to monetize a GenAI system when users privately differ in their preference for latency. Despite the apparent complexity of “customizing models,” the platform should train a single model that is aligned with user preference and screen with token caps. The theory provides a compact rationale for current industry practice and yields testable predictions about model training and token pricing.

There are multiple directions for future research to explore. First, incorporating the user prompting behavior would make the framework a better characterization of real-life usage of GenAI models. Second, given the fierce competition in the current AI industry, a competitive screening model is desirable for understanding the implication of having multiple platforms.

REFERENCES

- AMODEI, D., C. OLAH, J. STEINHARDT, P. CHRISTIANO, J. SCHULMAN, AND D. MANÉ (2016): “Concrete Problems in AI Safety,” *arXiv preprint arXiv:1606.06565*.
- ARMSTRONG, M. (1996): “Multiproduct Nonlinear Pricing,” *Econometrica*, 64, 51–75.
- ARMSTRONG, M. AND J. VICKERS (2001): “Competitive Price Discrimination,” *RAND Journal of Economics*, 32, 579–605.
- AZAGRA, D. AND D. STOLYAROV (2023): “Inner and outer smooth approximation of convex hypersurfaces. When is it possible?” *Nonlinear Analysis*, 230, 113225.
- BAI, Y., S. KADAVATH, S. KUNDU, ET AL. (2022): “Constitutional AI: Harmlessness from AI Feedback,” *arXiv preprint arXiv:2212.08073*.
- BERGEMANN, D., A. BONATTI, AND A. SMOLIN (2025): “The Economics of Large Language Models: Token Allocation, Fine-Tuning, and Optimal Pricing,” *ArXiv preprint arXiv:2502.07736*. URL: <https://arxiv.org/abs/2502.07736>.
- BLOEDEL, A. W. AND W. ZHONG (2024): “The cost of optimally-acquired information,” .
- BOLTON, P. AND M. DEWATRIPONT (2005): *Contract Theory*, Cambridge, MA: MIT Press.
- BOLTON, P. AND C. HARRIS (1999): “Strategic experimentation,” *Econometrica*, 67, 349–374.
- CAPLIN, A. AND M. DEAN (2015): “Revealed Preference, Rational Inattention, and Costly Information Acquisition,” *American Economic Review*, 105, 2183–2203.
- CHE, Y.-K., K. KIM, AND K. MIERENDORFF (2023): “Keeping the Listener Engaged: A Dynamic Model of Bayesian Persuasion,” *Journal of Political Economy*, 131, 3539–3573.
- CHE, Y.-K. AND K. MIERENDORFF (2019): “Optimal Dynamic Allocation of Attention,” *American Economic Review*, 109, 2993–3029.
- CHEN, D. AND W. ZHONG (2025): “Information Acquisition and Time-Risk Preference,” *American Economic Review: Insights*, 7, 213–230.
- FUDENBERG, D., P. STRACK, AND T. STRZALECKI (2018): “Speed, Accuracy, and the Optimal Timing of Choices,” *American Economic Review*, 108, 3651–3684.
- GUESS, A. M., N. MALHOTRA, J. PAN, P. BARBERÁ, H. ALLCOTT, T. BROWN, A. CRESPO-TENORIO, D. DIMMERY, D. FREELON, M. GENTZKOW, S. GONZÁLEZ-BAILÓN, E. KENNEDY, Y. M. KIM, D. LAZER, D. MOEHLER, B. NYHAN, C. V. RIVERA,

- J. SETTLE, D. R. THOMAS, E. THORSON, R. TROMBLE, A. WILKINS, M. WOJCIESZAK, B. XIONG, C. K. DE JONGE, A. FRANCO, W. MASON, N. J. STROUD, AND J. A. TUCKER (2023): “How do social media feed algorithms affect attitudes and behavior in an election campaign?” *Science*, 381, 398–404.
- HOFFMANN, J., S. BORGEAUD, A. MENSCH, ET AL. (2022): “Training Compute-Optimal Large Language Models,” *arXiv preprint arXiv:2203.15556*.
- KAMENICA, E. AND M. GENTZKOW (2011): “Bayesian Persuasion,” *American Economic Review*, 101, 2590–2615.
- KAPLAN, J., S. McCANDLISH, T. HENIGHAN, T. B. BROWN, B. CHESS, R. CHILD, S. GRAY, A. RADFORD, J. WU, AND D. AMODEI (2020): “Scaling Laws for Neural Language Models,” *arXiv preprint arXiv:2001.08361*.
- LAFFONT, J.-J. AND D. MARTIMORT (2002): *The Theory of Incentives: The Principal-Agent Model*, Princeton, NJ: Princeton University Press.
- LIN, Y., H. LIN, W. XIONG, S. DIAO, J. LIU, J. ZHANG, R. PAN, H. WANG, W. HU, H. ZHANG, H. DONG, R. PI, H. ZHAO, N. JIANG, H. JI, Y. YAO, AND T. ZHANG (2024): “Mitigating the Alignment Tax of RLHF,” in *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- MASKIN, E. AND J. RILEY (1984): “Monopoly with Incomplete Information,” *RAND Journal of Economics*, 15, 171–196.
- MATĚJKA, F. AND A. MCKAY (2015): “Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model,” *American Economic Review*, 105, 272–298.
- MILLI, S., M. CARROLL, Y. WANG, S. PANDEY, S. ZHAO, AND A. D. DRAGAN (2025): “Engagement, user satisfaction, and the amplification of divisive content on social media,” *PNAS Nexus*, 4, pgaf062.
- MORRIS, S. AND P. STRACK (2019): “The Wald Problem and the Relation of Sequential Sampling and Ex-Ante Information Costs,” Working paper. URL: https://scholar.princeton.edu/sites/default/files/smorris/files/wald_problem.ms.pdf.
- MUSSA, M. AND S. ROSEN (1978): “Monopoly and Product Quality,” *Journal of Economic Theory*, 18, 301–317.
- MYERSON, R. B. (1981): “Optimal Auction Design,” *Mathematics of Operations Research*, 6, 58–73.
- OUYANG, L., J. WU, X. JIANG, ET AL. (2022): “Training Language Models to Follow

- Instructions with Human Feedback,” *arXiv preprint arXiv:2203.02155*.
- ROCHET, J.-C. AND P. CHONÉ (1998): “Ironing, Sweeping, and Multidimensional Screening,” *Econometrica*, 66, 783–826.
- SANNIKOV, Y. AND W. ZHONG (2024): “Exploration and Stopping,” Working paper. URL: https://www.wjzhong.com/workingpapers/XS/Exploration_Stopping.pdf.
- SIMS, C. A. (2003): “Implications of Rational Inattention,” *Journal of Monetary Economics*, 50, 665–690.
- VASWANI, A., N. SHAZEER, N. PARMAR, J. USZKOREIT, L. JONES, A. N. GOMEZ, Ł. KAISER, AND I. POLOSUKHIN (2017): “Attention is all you need,” *Advances in neural information processing systems*, 30.
- ZHONG, W. (2022): “Optimal dynamic information acquisition,” *Econometrica*, 90, 1537–1582.