

Subsidizing Sequential Search

[Most recent version: [link](#)]

Salvador Candelas* Nicole Immorlica[†] Brendan Lucier[‡]

December 20, 2025

Abstract

We study markets where firms compete for consumer attention by subsidizing costly product inspection. These subsidies do not change product quality, but they alter the order in which consumers search by lowering inspection costs. We establish a subsidy–sorting principle: in any equilibrium, higher–quality firms provide weakly larger subsidies, leading consumers to search in descending subsidy order. A unique equilibrium survives forward–induction reasoning in the spirit of the Intuitive Criterion: low–quality firms are never inspected, intermediate–quality firms separate with strictly increasing subsidies, and high–quality firms pool at the full subsidy. This equilibrium maximizes information revelation among all possible outcomes and ensures efficient inspection. We then extend the analysis to AI–mediated platforms that can create and price “inspection tokens.” The platform’s optimal linear pricing leads to excessive inspection relative to the social optimum. While this distortion does not reduce consumer welfare, it reallocates surplus from sellers to the platform.

Keywords: Sequential Search, Subsidy Competition, Attention Allocation, Agentic Economies

JEL Codes: D47, D83, L86, M37

*Pennsylvania State University. Email: scandelas@psu.edu

[†]Microsoft Research / Yale. Email: nicimm@gmail.com

[‡]Microsoft Research. Email: brlucier@microsoft.com

1 Introduction

Markets run on attention. Buyers want to find a product that fits their needs, but evaluating options takes time and effort. This challenge is familiar in online markets: a consumer types a query, browses links, and exerts effort before finding what she wants. Firms compete for the consumer’s attention not only through price and product features but also by lowering the cost of evaluation, most visibly through sponsored placements and promoted listings. This phenomenon arises also beyond the web. For instance, graduate programs cover travel expenses so prospective students can visit campus and meet faculty; car dealers hand out cash for test drives; galleries and non-profits host receptions to draw people. Across these settings, in both online and off-line markets, the consumer faces search frictions and the other side competes to offset these search frictions.

We formalize this problem as a new kind of “directed search”: a consumer chooses which options to evaluate, and each evaluation carries a cost. Firms compete for her attention by subsidizing some or all of her search cost. These subsidies play a dual role. They are instrumental, because they reduce the cost of being considered and enlarge the set of products the consumer is willing to examine. They are informative, because only firms confident in high match value find it worthwhile to subsidize. In this way, subsidies both broaden what is searched and determine the order in which search unfolds.

We formalize these ideas in [Section 2](#). We study a sequential-search model with vertical match uncertainty. A consumer seeks a product from one of n firms. Firm j has a privately known type $t_j \in [0, 1]$, which denotes the probability that its product is a match for the consumer’s needs. Firm types are drawn i.i.d. from a common prior. To learn whether a firm’s product matches, the consumer must “inspect” it at cost $c > 0$, which is a commonly known search or inspection cost. Each firm j , after observing its type, simultaneously announces a subsidy $s_j \in [0, c]$ paid to the consumer upon inspection. The consumer observes the subsidy profile $\mathbf{s} = (s_1, \dots, s_n)$, updates beliefs about match probabilities, and chooses an inspection order that maximizes expected utility, trading off lower net costs $(c - s_j)$ against inferred quality. We analyze the Perfect Bayesian equilibria of this game and show how subsidized inspection reallocates attention in equilibrium.

Before we describe the analysis of this setting, we relate it to two well-known studies. One way to view our analysis is as a Pandora’s Box problem in the spirit of [Weitzman \(1979\)](#) in which the “boxes” can play an active role by subsidizing search. In this spirit, the model captures decentralized search in which privately-informed firms try to attract consumers by reducing their inspection costs. A second way to view our analysis is as a platform in the spirit of [Athey and Ellison \(2011\)](#) in which firms go beyond paying for prominence, and instead can subsidize search. After describing our analysis, we highlight how this approach

to search on platforms is relevant to the emerging trend of *agentic search*, where consumers delegate discovery and evaluation to AI assistants.¹

Three questions guide our investigation of this directed-search framework. First, how do subsidies signal quality? Second, in this signaling game, do forward induction refinements in the spirit of [Cho and Kreps \(1987\)](#) refine equilibria? Third, how do market fundamentals, competition, search costs, and subsidy prices, affect consumer and social welfare? How do changes in market fundamentals, more competition, lower search costs, higher subsidy prices, affect consumer and social welfare?

[Section 3](#) addresses the first question. We establish a *subsidy-sorting principle*: in every equilibrium, higher equilibrium levels of subsidy are rewarded with higher attention, which implies that higher-type firms offer (weakly) higher subsidies. Therefore, subsidies on the equilibrium path are a positive signal of quality.² Given the subsidy-sorting principle, the consumer’s optimal search strategy is therefore to inspect firms in descending order of subsidy, breaking ties randomly. Search stops when a match is found or when the next inspection yields a negative expected payoff.

We next address equilibrium multiplicity and ask which outcomes survive a natural forward-induction refinement. Following [Cho and Kreps \(1987\)](#) and [Banks and Sobel \(1987\)](#), we impose a forward-induction restriction that eliminates equilibria supported by implausible off-path beliefs. Our main result establishes that this refinement selects a unique equilibrium outcome in which every firm type that can profitably induce inspection is inspected. This equilibrium outcome is supported by a distinctive step-increasing-step subsidy pattern.³




Figure 1: Low types are never inspected, intermediate types separate with strictly increasing subsidies, and high types pool at the full subsidy.

¹Examples include interfaces such as *ChatGPT*, *Gemini*, and other LLM-based chatbots, as well as Shopify’s *Sidekick*, Expedia’s conversational trip-planner, and Amazon’s *Rufus*.

²In some equilibria, higher *off-path* subsidies could convey a negative signal of value.

³Our description below and [Figure 1](#) uses a subsidy schedule in which all firms whose types are below $\bar{t}$ offer zero subsidies. There exist outcome-equivalent equilibrium that satisfies the refinement in which these firms offer strictly positive subsidies that are nevertheless insufficient to attract the consumer’s attention.

We illustrate this pattern in [Figure 1](#). Low-quality firms offer zero subsidy and are never inspected. Intermediate types separate by offering strictly increasing subsidies, thus revealing their quality. High types pool at the full subsidy, reimbursing the inspection cost entirely. Pooling generates a fundamental tension: it obscures differences among the best firms, yet it makes their inspection free for consumers. The consumer gains fully from this trade-off, but social surplus suffers from the inefficient search due to the loss of information at the top.

[Section 5](#) addresses the welfare question: how do changes in market primitives affect consumer and social outcomes? As expected, more competing firms or lower search costs increase consumer welfare. Higher per-unit subsidy prices reduce expected transfers and discourage firm participation, both lowering consumer surplus.

We apply these insights to the emerging *agentic economy* ([Rothschild et al., 2025](#)). AI assistants are reshaping the economics of consumer search.⁴ Unlike human shoppers who rely on salience and limited patience, algorithmic agents follow programmed utility functions and face measurable costs. Each evaluation is a sequence of auditable computational events (e.g., API calls, database queries, or token usage). This transparency creates new levers for market manipulation: sellers can now purchase algorithmic attention directly, subsidizing inspection to push their products ahead of competitors.

[Section 6](#) addresses the platform question: how should a monopolistic platform price computational attention, and what are the welfare consequences? The platform sets token prices that determine which firms can bid for visibility and how much consumers ultimately search. We show that the platform maximizes profits when search is inefficient. By lowering token prices, it induces pooling among top firms, expands entry, and forces agents to inspect more options than is socially optimal. At the profit-maximizing price, the additional attention and search more than offset the lower per-inspection charge. This strategy raises revenue while appearing to benefit consumers through cheaper individual inspections. In equilibrium, the platform exploits the wedge between consumer welfare and social surplus, extracting rents from the opacity it creates.

[Section 7](#) concludes. All omitted proofs are in Appendices. The remainder of this introduction places our contribution within the context of the existing literature.

Relation to the literature. The classical search literature treats inspection as a fixed burden borne by consumers. Early work frames search as an optimal stopping problem ([Stigler, 1961](#); [McCall, 1970](#); [Diamond, 1971](#)). [Weitzman \(1979\)](#) provides the canonical index rule when inspection costs are heterogeneous. Subsequent work extends this framework to directed and simultaneous search under match-quality uncertainty ([Chade and Smith, 2006](#)).

⁴For example, a shopper might instruct an AI assistant to “scan every retailer, check return policies, and buy the cheapest 32-inch monitor that ships by Friday”.

We build on this tradition but depart by making search costs contractible: firms can subsidize evaluation, turning a stopping problem into an equilibrium allocation of attention.

A growing strand studies strategic information provision and disclosure in search environments. Firms strategically disclose more when they mainly compete to attract new customers and cannot condition on buyer beliefs; they disclose less when competition is about persuading already-arrived buyers or when firms can condition on beliefs (Board and Lu, 2018; Au and Whitmeyer, 2023). Recent work recasting Weitzman’s Pandora’s box in competitive and intermediary settings finds that competition often improves informativeness, yet principals who sell signals can implement efficient stopping while extracting surplus (Ding et al., 2025; Mekonnen et al., 2025). These contributions show that strategic information provision reshapes inspection order, and thus welfare. Our analysis complements this literature but shifts the instrument from beliefs to search frictions: rather than solely manipulating posteriors, firms in our model manipulate the cost of learning. Subsidies lower frictions directly and thereby reorder the search path, benefiting the consumer.

A second strand examines how firms compete for consumer attention through advertising (Anderson and Renault, 2006; Edelman et al., 2007; Athey and Ellison, 2011; Chen and He, 2011). In sponsored search, firms bid for ad positions, with higher bidders receiving more prominent placement. The generalized second-price auction creates efficient rankings where firms with higher valuations for clicks secure better positions (Edelman et al., 2007). Athey and Ellison (2011) show this extends to organic search when firms can influence their ranking through quality investments. We demonstrate that subsidizing inspection achieves similar efficient sorting without auctions or slot constraints. The mechanism is simpler and more direct. Yet we uncover a crucial difference: subsidy caps create pooling at the top, generating excessive search that hurts social welfare while preserving consumer surplus. This distinguishes our setting from position auctions where firms can always separate through higher bids.

The obfuscation literature highlights the opposite strategic move: firms or platforms can raise search frictions to soften competition and extract rents (Ellison and Wolitzky, 2012; Nocke and Rey, 2024), and recent work shows how algorithmic obfuscation and sponsored slots can be used jointly to boost platform revenue (Janssen et al., 2025). We show how those two levers can be combined: a platform can partly obfuscates organic listings by creating stronger competition among firms. This competition leads to lower search frictions as firms subsidize inspection more aggressively and hit subsidy caps. Consumers perceive inspections as effectively free and search more, while the platform extracts revenue (from firms) through greater search volume. This mechanism nests cleanly in the two-sided platform literature: subsidy pricing is both a participation instrument and an attention-pricing device (Rochet and Tirole, 2003; Weyl, 2010).

2 Environment

A consumer seeks a product that matches her needs. The benefit that she accrues from finding such a product is $u > 0$; otherwise, she obtains her outside option, $\emptyset$, which delivers zero utility. She searches for the product on a platform, inspecting the products of different firms. Each inspection costs $c > 0$ and reveals whether the inspected firm's product meets her needs.⁵ The consumer searches sequentially and terminates either when she finds a match or she opts out.⁶

Each firm $j \in \mathcal{J} = \{1, \dots, n\}$ has a type $t_j \in T := [0, 1]$, drawn independently from a common distribution F with full support.⁷ The interpretation is that t_j is the firm's match probability upon inspection. Hence, we refer to a higher type firm as *better* and a lower type firm as *worse*. A firm earns $v = 1$ if its product matches the consumer, and 0 otherwise. Before inspections occur, a firm may offer a subsidy $s_j \in [0, c]$ that is paid if the consumer inspects the firm's product. The subsidy reduces the consumer's effective inspection cost from c to $c - s_j$. Each unit of subsidy costs the firm $p \in \mathbb{R}_+$ per unit. In [Section 6](#), we consider a profit-maximizing choice of p for a platform, but for the baseline model, p is a primitive.

Timing. Each firm privately observes t_j and simultaneously chooses a subsidy $s_j \in [0, c]$. After observing the vector of subsidies $\mathbf{s} = (s_j)_{j \in \mathcal{J}}$, the consumer selects a sequential inspection rule. Upon inspecting firm j , the consumer learns whether the product satisfies her need (i.e., whether a match occurs).

2.1 Strategies

Each per-unit price $p \in \mathbb{R}_+$ induces a Bayesian game $\mathcal{G}(p)$ between the consumer and the set of firms $\mathcal{J} = \{1, \dots, n\}$.

A *subsidy policy* is a measurable function $\sigma : T \rightarrow [0, c]$, mapping a type t to a subsidy $s = \sigma(t)$. A subsidy policy profile $(\sigma_j)_{j \in \mathcal{J}}$ is a collection of subsidy policy stating each firm j policy. Given a type realization $\mathbf{t}$, a subsidy policy profile $(\sigma_j)_{j \in \mathcal{J}}$ induces a realized subsidy profile $\mathbf{s} = (s_j)_{j \in \mathcal{J}}$ with $s_j = \sigma_j(t_j)$.

The consumer observes $\mathbf{s}$ and then implements an inspection rule. An inspection rule $r \in \mathcal{R}$ is a planned order of inspections with an outside-option stop: inspections proceed sequentially until a match is found or the outside option is reached. Formally, let $\mathcal{R}$ be the set of permutations of $\tilde{\mathcal{J}} = \mathcal{J} \cup \{\emptyset\}$. For $r \in \mathcal{R}$, write $i \succ_r j$ to mean " i is scheduled before j ,"

⁵We assume that the inspection cost is not so large that all search is deterred. Formally, $c < (1 + pu)/p$, so that some firm types can profitably subsidize inspection and induce the consumer to search.

⁶We focus on inspection rules that stop either after all planned inspections fail or whenever a match is formed. This is a minimum requirement for optimality in this search environment.

⁷The full-support assumption is not crucial, but simplifies the exposition.

i.e., $r(i) > r(j)$; by convention, if $\emptyset \succ_r j$ then j is never inspected. The consumer's strategy is an *inspection policy*, a stochastic kernel $\iota : \mathcal{R} \times [0, c]^n \rightarrow [0, 1]$, where $\iota(r \mid \mathbf{s})$ denotes the probability that inspection rule r is used at subsidy profile $\mathbf{s}$.

The consumer's *beliefs* are given by a measurable mapping $\mu : [0, c]^n \rightarrow \Delta(T)^n$ that assigns to each subsidy profile $\mathbf{s}$ a posterior over types; our formulation assumes both on- and off-path that the consumer's beliefs about a firm j do not depend on the subsidies offered by other firms.⁸

An *assessment* is a triple $((\sigma_j)_{j \in \mathcal{J}}, \iota, \mu)$, where σ_j is firm j 's subsidy policy, ι is the consumer's inspection policy, and μ is the consumer's belief system.

Throughout, "increasing" means weakly increasing, i.e. $x' \geq x \implies f(x') \geq f(x)$. We use "strictly increasing" when the inequality is strict. Analogous definitions apply for decreasing functions and monotone functions.

2.2 Equilibrium Concept

Our solution concept is Perfect Bayesian Equilibrium (PBE). An assessment $((\sigma_j)_{j \in \mathcal{J}}, \iota, \mu)$ is a PBE if (i) each firm's strategy maximizes expected profit, (ii) the consumer's strategy maximizes expected utility given beliefs, and (iii) beliefs are updated by Bayes' rule whenever possible. We restrict attention to symmetric equilibria.

In a symmetric Perfect Bayesian Equilibrium, two restrictions apply. First, *exchangeability*: the consumer treats firms offering identical subsidies identically, both on and off the equilibrium path.⁹ This implies that the belief system can be characterized by a measurable mapping $\mu : [0, c] \rightarrow \Delta(T)$ to which we refer as the consumer's belief function.¹⁰ Second, all firms adopt the same subsidy policy σ . Thus a symmetric assessment can be written as (σ, ι, μ) , where σ is the common subsidy policy, ι the inspection policy, and μ the consumer's belief function. Henceforth, we refer to symmetric PBE as *equilibria*.

A key object is the *attention function*. A profile (σ, ι) induces a distribution over inspection paths and outcomes. Offering a subsidy s delivers a level of attention, defined as the ex ante probability that the firm is inspected. Any given firm is inspected if, and only if, all firms scheduled before it fail to match. To encode the outside option, define $t_\emptyset := 1$. The attention

⁸Given that firms' types are independent, the stochastic independence of posterior beliefs would be implied by [Kreps and Wilson \(1982\)](#) (were the game to have finitely many actions).

⁹Formally, for any two subsidy profiles $\mathbf{s}$ and $\mathbf{s}'$, if there exists a permutation π such that $s_j = s'_{\pi(j)}$ for all $j \in \mathcal{J}$, then $\iota(\mathbf{s}) = \iota(\mathbf{s}')$ and $\mu(\mathbf{s}) = \mu(\mathbf{s}')$.

¹⁰Formally, under symmetry and stochastic independence of posteriors, the conditional distribution of a firm's type given that it offers subsidy s is identical across firms and independent of other firms' strategies.

function of a firm offering s under (σ, ι) is

$$q_{\sigma, \iota}(s) := \int \left(\sum_{r \in \mathcal{R}} \iota(r \mid (s, \sigma(\mathbf{t}_{-j}))) \prod_{i \in \tilde{\mathcal{J}}: j <_r i} (1 - t_i) \right) dF^{n-1}(\mathbf{t}_{-j}). \quad (1)$$

In words, the *attention function* at s , $q_{\sigma, \iota}(s)$, is the ex ante probability that a firm offering subsidy s is inspected under equilibrium play.

From an interim perspective, a type- t firm that offers subsidy s is inspected with probability $q_{\sigma, \iota}(s)$, succeeds in matching with probability t , and pays ps upon inspection. The firm's *interim* profit is therefore

$$\pi_{\sigma, \iota}(s; t) = (t - ps) q_{\sigma, \iota}(s). \quad (2)$$

Turning to the consumer, a match fails to occur under rule r precisely when every firm scheduled before $\emptyset$ fails to match. Thus, given a posterior belief $\mu(\mathbf{s})$ over types, the consumer expects to match on expectation $\mathbb{E}_{\mathbf{t} \sim \mu(\mathbf{s})} \left[1 - \prod_{i \in \tilde{\mathcal{J}}: i <_r \emptyset} (1 - t_i) \right]$.

Definition 1. An equilibrium of the game $\mathcal{G}(p)$ is a triple (σ, ι, μ) such that:

(i) **Firm optimality.** Given (σ, ι) , each type- t firm chooses

$$\sigma(t) \in \arg \max_{s \in [0, c]} (t - ps) q_{\sigma, \iota}(s).$$

(ii) **Consumer optimality.** Given a realized subsidy profile $\mathbf{s}$ and posterior beliefs $\mu(\mathbf{s})$, the inspection distribution $\iota(\mathbf{s}) \in \Delta(\mathcal{R})$ assigns positive probability only to rules that maximize expected utility:

$$\text{supp}(\iota(\mathbf{s})) \subseteq \arg \max_{r \in \mathcal{R}} \mathbb{E}_{\mathbf{t} \sim \mu(\mathbf{s})} \left[u \left(1 - \prod_{i \in \tilde{\mathcal{J}}: i <_r \emptyset} (1 - t_i) \right) - \sum_{j \in \tilde{\mathcal{J}}} (c - s_j) \prod_{i \in \tilde{\mathcal{J}}: i <_r j} (1 - t_i) \right].$$

(iii) **Belief consistency.** μ is derived from σ by Bayes' rule wherever possible.

Equilibria always exist. For example, if all firms set zero subsidies, the consumer searches randomly on the equilibrium path and assigns off-path beliefs that any subsidizing firm is of type 0, no such firm is inspected. In [Section 4](#) we construct a consumer-optimal equilibrium, which applies a forward-induction logic in the spirit of [Cho and Kreps \(1987\)](#) and rules out reliance on implausible off-path beliefs.

2.3 Equilibrium Refinement

Following [Cho and Kreps \(1987\)](#) and [Banks and Sobel \(1987\)](#), we refine equilibria at off-path beliefs. We require that at off-path subsidy s the consumer do not place positive probability in *dominated* types. The idea is as follows: after an off-path subsidy, the consumer assigns zero probability to type t if there exists some other type t' for whom that subsidy yield a higher payoff than the the equilibrium payoff under any posterior belief and its corresponding inspection policy. Thus, when evaluating off-path moves, the consumer rules out firm types for whom that deviation is not dominant.

Fix an equilibrium (σ, ι, μ) of $\mathcal{G}(p)$. A subsidy $s \in [0, c]$ is *off-path* if there is not type- t that offers it, that is, $s \notin \text{Im}(\sigma)$. Let $\mathcal{C}(\sigma)$ denote the set of admissible consumer's continuations based on σ : each element is a pair (ι', μ') such that (i) the consumer's inspection policy ι' is optimal given beliefs μ' , (ii) exchangeability holds, and (iii) on-path beliefs obey Bayes' rule. Given $(\iota', \mu') \in \mathcal{C}(\sigma)$, the expected profit for a type- t firm that offers s is $\pi_{\sigma, \iota'}(s; t) = (t - ps) q_{\sigma, \iota'}(s)$.

Definition 2. For an off-path subsidy level s , we say that (t, s) is dominated by the equilibrium value if

$$\pi_{\sigma, \iota}(\sigma(t); t) > \max_{(\iota', \mu') \in \mathcal{C}(s)} \pi_{\sigma, \iota'}(s; t).$$

That is, type- t cannot improve its equilibrium payoff by deviating to the off-path subsidy s under any admissible continuation.

Moreover, we say that type t is dominated by type t' at s if:

$$\pi_{\sigma, \iota}(\sigma(t); t) \leq \pi_{\sigma, \iota'}(s; t) \implies \pi_{\sigma, \iota}(\sigma(t'); t') < \pi_{\sigma, \iota'}(s; t') \quad \forall (\iota', \mu') \in \mathcal{C}(\sigma)$$

Let $t \in D_{\sigma, \iota}(s)$ if (t, s) is dominated by the equilibrium value or there exists some t' such that t is dominated by t' at s .

Definition 3. An equilibrium (σ, ι, μ) fails the equilibrium refinement if, there exists some off-path subsidy s , and a type t such that $t \in \text{supp}(\mu(s))$ and $t \in D_{\sigma, \iota}(s)$.¹¹

Operationally, equilibrium-dominance restricts off-path beliefs to the *economically credible* types for whom the deviation could be optimal. In [Section 4](#), we show that there is a unique equilibrium outcome that survives this refinement and characterize the set of equilibria that achieves it.

¹¹Abusing notation, we write $\mu(s)$ for the marginal distribution of the first coordinate of μ evaluated at s . By symmetry this marginal is the same for every firm, so we omit any firm index.

3 The Subsidy–Sorting Principle

Do high–quality firms subsidize more than low–quality ones, or can competition overturn this relationship? This section establishes the *Subsidy–Sorting Principle*. In every equilibrium, the attention a firm receives is increasing in its subsidy. This monotonicity implies that equilibrium subsidy policies σ are increasing in type, and that the consumer follows a descending–subsidy inspection rule: firms are visited in descending order of their realized subsidies, with uniform tie–breaking within pools. Search terminates either upon a successful match or when the next inspection yields negative expected surplus.

Theorem 1 (Subsidy–Sorting Principle). *In every PBE of the game $\mathcal{G}(p)$, a higher subsidy generates more attention. Moreover,*

1. **Monotone subsidies:** *the subsidy policy is increasing in type, i.e., if $t' \geq t$ then $\sigma(t') \geq \sigma(t)$.*
2. **Descending–subsidy index rule (DSIR):** *the consumer inspects firms in descending order of their realized subsidies, breaking ties uniformly. After a failed inspection, she continues search only if for the next inspected subsidy level s the expected benefit is higher than the expected cost, i.e. $\mathbb{E}[t \mid \sigma(t) = s] \geq \frac{c-s}{u}$.*

The above result offers several structural properties that must be satisfied by any equilibrium. Higher subsidies draw more inspections. Better types exploit this by offering larger subsidies. The consumer then inspects in descending order of subsidies, breaking ties at random, and stops once further search is not worthwhile.

The reasoning unfolds in two steps. First, inspection probabilities are increasing in subsidy, so subsidies sort firms by how much attention they receive. Second, because firms' profits have a single-crossing property, higher types benefit more from extra attention and therefore choose higher subsidies. Once subsidies are monotone, descending-subsidy search is optimal: higher subsidies both lower net costs and signal higher chances of a match, and the consumer stops when the expected gain turns negative.

3.1 Subsidy Monotonicity

Fix an equilibrium. We argue that offering a higher subsidy must increase the (ex-ante) inspection probability. Consider two on-path subsidies $s < s'$, i.e., there exists t and t' such that $\sigma(t) = s$ and $\sigma(t') = s'$. We claim that $q(s')$, the ex-ante probability that the consumer inspects given subsidy s' , is higher than $q(s)$, the corresponding probability given subsidy s . Towards a contradiction, suppose that $q(s') < q(s)$. Then, type t' payoff from deviating to a

subsidy s would be strictly higher than its equilibrium payoff:

$$(t' - ps)q(s) > (t' - ps)q(s') \geq (t' - ps')q(s'),$$

where the first inequality follows from $q(s) > q(s')$, and the second from $s < s'$.¹² Therefore, type t' would have a strictly profitable deviation. Putting it simply, if a firm sacrificed attention by subsidizing search, then it would have no interest in doing so.

The second step uses monotone comparative statics to show that better firms have a greater interest in subsidizing search. As attention is more rewarding for higher types, they have a stronger interest in obtaining it. Specifically, each firm's payoff over inspection probabilities satisfies single-crossing differences (Milgrom and Shannon, 1994) in type. Suppose that type t (resp., t') offers subsidy s (resp., s') and obtains an inspection probability q (resp., q'). We argue that if $t' \geq t$, then $q' \geq q$ by evaluating each type's best response. In equilibrium, type t' must favor subsidy s' over s , which implies

$$(t' - ps')q' \geq (t' - ps)q.$$

Analogously, type t must favor subsidy s over s' , which implies

$$(t - ps)q \geq (t - ps')q'.$$

Rearranging the above inequalities yields

$$t'(q' - q) \geq p(q's' - qs) \geq t(q' - q).$$

Now, if it were the case that $q' < q$, the above inequality would imply that $t' < t$, which results in a contradiction.

Pasting these two steps together, we obtain that: (i) the consumer pays more attention to firms that offer higher subsidies, and (ii) better firms have more to gain from the consumer's attention. This establishes that, in every equilibrium, higher types must offer greater subsidies and be rewarded with more attention.

Lemma 1. *In any equilibrium (σ, ι, ρ) , the inspection probability $q_{\sigma, \iota}$ is increasing in on-path subsidies, and the subsidy policy σ is increasing in type.*

¹²The second inequality is not strict as nothing prevents $q(s') = 0$.

3.2 Consumer Search: Ordering and Stopping

Having established that equilibrium subsidies are weakly increasing in type, we now turn to the consumer's problem. Higher types offer higher subsidies and, as a result, receive more attention. This suggests that, in expectation, firms with larger subsidies should be inspected earlier. We show that this intuition is correct: the consumer's optimal policy is to inspect in descending order of a simple index that is monotone in subsidies on the equilibrium path.

Given a subsidy policy σ and a belief system $\mu : [0, c] \rightarrow \Delta(T)$, a realized subsidy s reveals that the issuing firm's type lies in the level set $\sigma^{-1}(s)$. Because the consumer's payoff is linear in the match indicator, all payoff-relevant information is summarized by the *posterior success probability*

$$\tau_\mu(s) := \mathbb{E}_{t \sim \mu(s)}[t]. \quad (3)$$

When $s \in S_\sigma := \{\sigma(t) : t\}$ and μ is Bayes consistent, $\tau_\mu(s) = \mathbb{E}[t \mid \sigma(t) = s]$. If $s \notin \text{Im}(\sigma)$, $\tau_\mu(s)$ is determined by off-path beliefs.

Fix a realized subsidy profile $\mathbf{s} = (s_j)_{j \in \mathcal{J}}$ with $s_j \in S_\sigma$ for all j . Write $\tau_j := \tau_\mu(s_j)$ and $\kappa_j := c - s_j$ for the posterior match probability and the net inspection cost. Inspecting firm j yields a match with probability τ_j and costs κ_j . The expected payoff is therefore $\tau_j(u - \kappa_j) + (1 - \tau_j)(V - \kappa_j)$, while stopping delivers V . Inspection is optimal when the former exceeds V , which defines the reservation index

$$r(\tau, \kappa) := u - \frac{\kappa}{\tau}, \quad (4)$$

with the convention $r(\tau, \kappa) = -\infty$ when $\tau = 0$. Inspection is optimal if and only if $V < r(\tau, \kappa)$. This is exactly the index rule in [Weitzman \(1979\)](#) for a Bernoulli prize with success probability τ and cost κ .

We introduce the *descending-subsidy index rule* under beliefs μ .

Definition 4 (Descending-subsidy index rule). Fix $\mu : [0, c] \rightarrow \Delta(T)$ and define

$$r_\mu(s) := u - \frac{c - s}{\tau_\mu(s)}, \quad (5)$$

with $r_\mu(s) = -\infty$ if $\tau_\mu(s) = 0$. The descending-subsidy index rule ($DSIR_\mu$) prescribes:

- (a) **Ordering:** inspect firms in descending order of $r_\mu(s)$, breaking ties uniformly at random;
 - (b) **Stopping:** terminate upon the first match, or whenever the next available index is nonpositive.
- The search policy that implements $DSIR_\mu$ for any realized subsidy profile is the descending-subsidy index policy, $DSIP_\mu$.

The next result states that the DSIP is consumer optimal. Moreover, it is unique under exchangeability. Note that the result does not require $S_\sigma = [0, c]$. If a realized subsidy $s \notin S_\sigma$

is observed, $DSIR_\mu$ remains well defined once $\tau_\mu(s)$ is specified by off-path beliefs; different off-path beliefs can support different equilibria.

Lemma 2. *For any belief system μ , $DSIP_\mu$ maximizes the consumer's expected utility. It is unique up to tie-breaking among firms with equal index.*

Based on this insights, we can easily describe the consumer search patterns on equilibrium. Whenever the subsidy policy is increasing, the consumer should always prefer inspecting a higher-subsidy firm: it offers both a higher probability of match and a lower net cost. If the consumer did not inspect such a firm before a lower-subsidy one, she could strictly improve her payoff by reordering. Once a match is found, she has no incentive to continue, since further inspections only add cost without raising her payoffs. And if no match has been found, inspection should stop once the next available firm yields negative expected surplus, since the outside option guarantees a payoff of zero. This logic pins down the consumer's optimal rule: inspect in descending order of subsidies, break ties at random, stop at the first match, and otherwise continue only while the expected net benefit of the next inspection is positive.

Corollary 1. *Suppose σ is weakly increasing and μ is Bayes consistent on S_σ . Then $r_\mu(\cdot)$ is increasing on S_σ . Consequently, $DSIR_\mu$ reduces to inspecting firms in weakly descending order of realized subsidies, breaking ties uniformly, and continuing search only while $\mathbb{E}[t \mid \sigma(t) = s] > \frac{c-s}{u}$.*

Combining [Lemma 1](#) with [Corollary 1](#) immediately yields the Subsidy-Sorting Principle: higher types offer weakly higher subsidies; the consumer inspects in descending subsidy order; and she stops exactly when she matches or when the next firm expected profits are lower than the net costs.

4 The Reasonable Equilibrium

We now characterize the unique equilibrium outcome that survives forward-induction refinement. We refer to the equilibrium that induces it the *reasonable equilibrium*. This equilibrium takes a simple shape: low types do not subsidize, intermediate types separate with a strictly increasing schedule, and high types pool at the cap c . This equilibrium induces efficient participation: every firm that could profitably signal so as to secure positive inspection does so and is inspected with positive probability. The formal result is stated in [Theorem 2](#).

Our construction proceeds in four steps. First, we derive firm incentives under full separation, yielding a differential equation that characterizes truth-telling subsidies. Second, we identify the lowest participating type, which pins down the admissible boundary condition. Third, if the implied policy exceeds the subsidy cap c , we iron back the schedule by pooling

higher types at the cap. Finally, we establish that this outcome is the only one consistent with refinement.

4.1 Firm Incentives under Full Separation

Assume a strictly increasing subsidy policy σ . Under full separation, subsidies reveal types one-to-one. By Lemma 2, the consumer should inspect firms in descending order of subsidies, breaking ties uniformly, and stops once the next inspection fails the reservation test. Let $\underline{t}$ denote the lowest type inspected.¹³ Firms with $t < \underline{t}$ are never inspected.

Consider the incentives of a type- t firm. If it mimics some type $\tilde{t}$, it offers subsidy $\sigma(\tilde{t})$ and is inspected with probability $q_n^{\text{sep}}(\tilde{t}; \underline{t})$. The key property of this inspection probability is that, under full separation, it depends only on $\underline{t}$ and not on the detailed shape of σ .¹⁴ The firm's expected payoff from mimicking $\tilde{t}$ is

$$\pi(\tilde{t}; t) = [t - p\sigma(\tilde{t})] q_n^{\text{sep}}(\tilde{t}; \underline{t}). \quad (6)$$

In equilibrium, truth-telling must maximize (6). Since σ is strictly increasing and hence almost everywhere differentiable, the envelope condition at $\tilde{t} = t > \underline{t}$ implies

$$-p\sigma'(t) q_n^{\text{sep}}(t; \underline{t}) + [t - p\sigma(t)] (q_n^{\text{sep}})'(t; \underline{t}) = 0,$$

or equivalently that marginal cost should be equal to marginal benefit:

$$p \frac{d}{dt} [\sigma(t) q_n^{\text{sep}}(t; \underline{t})] = t (q_n^{\text{sep}})'(t; \underline{t}). \quad (7)$$

Every strictly increasing, truth-telling subsidy policy must satisfy (7).

4.2 Boundary Condition from Participation and Inspect-Worthiness

To integrate (7), we require a boundary condition. The relevant support of the subsidy policy begins at the lowest type that both finds participation profitable and is just worth inspecting. These two primitives jointly determine the boundary.

A firm that expects to be inspected with positive probability participates only if it expects nonnegative profit, i.e. $\sigma(t) \leq t/p$. The consumer inspects only if the reservation index is nonnegative, which under separation requires $\sigma(t) \geq c - t$. Together these inequalities

¹³Under full separation, the subsidy perfectly reveals t , so the reservation index is $r(t) = u - (c - \sigma(t))/t$. Inspection requires $r(t) \geq 0$, hence the lowest inspected type solves $\sigma(\underline{t}) = c - u\underline{t}$. Whenever $\sigma(1) > c - u$, the solution exists and is unique, with $\underline{t} \in (0, c)$. Otherwise, no type is inspected under this policy, set $\underline{t} = 1$.

¹⁴A closed form is derived in Appendix C.1, where we show $q_n^{\text{sep}}(t; \underline{t}) = (1 - \int_{\underline{t}}^1 x dF(x))^{n-1} \mathbf{1}\{t \geq \underline{t}\}$.




Figure 2: The shaded regions mark subsidy levels that make trade infeasible. In red, inspection yields negative payoff for the consumer (net costs exceed expected gains). In blue, firms incur losses (expected sales do not cover the subsidy). The dot denotes the cutoff pair $(\underline{t}, \underline{s})$, the highest low type that leaves both sides exactly indifferent. In this figure: $c = 0.5$ and $p = 1$.

define a feasible corridor where trade is efficient:

$$\max\{0, c - t\} \leq \sigma(t) \leq \frac{t}{p}.$$

Inefficient trade means that some party, either the consumer or the firm, expected negative payoff of such transaction. Figure 2 illustrates such region.

The two bounds meet at a unique cutoff $\underline{t}$. This cutoff plays a dual role: it is the marginal participant, breaking even on subsidies and exactly at the threshold of inspect-worthiness:

$$\underline{t}(p, c) = \frac{pc}{1+p} \in (0, 1], \quad \underline{s} = \frac{\underline{t}(p, c)}{p} = \frac{c}{1+p}.$$


Solving Equation (7) with boundary condition $(\underline{t}, \underline{s})$ yields a strictly increasing, incentive-compatible subsidy policy for $t \in [\underline{t}, 1]$:

$$\sigma^{\text{sep}}(t; \underline{t}) = \frac{t}{p} - \frac{\int_{\underline{t}}^t q_n^{\text{sep}}(x; \underline{t}) dx}{p q_n^{\text{sep}}(t; \underline{t})}, \quad (8)$$

with $\sigma^{\text{sep}}(t; \underline{t}) = 0$ for $t < \underline{t}$. We refer to Equation (8) as the *separating policy* with support on $[\underline{t}, 1]$, since it arises under full separation when subsidies perfectly sort types.

4.3 Ironing Back to the Original Constraints

The revealing schedule $\sigma^{\text{sep}}(\cdot; \underline{t})$ solves the firm's problem under full separation but ignores the statutory cap $\sigma(t) \leq c$. Restoring the cap leaves two cases. If $\sigma^{\text{sep}}(1; \underline{t}) \leq c$, then $\sigma^{\text{sep}}(\cdot; \underline{t})$ is already feasible. If instead $\sigma^{\text{sep}}(1; \underline{t}) > c$, the upper tail must be truncated and high types pooled at the maximum subsidy $s = c$.



(a) Feasible without cap: if $\sigma^{\text{sep}}(1; \underline{t}) \leq c$, the revealing schedule requires no adjustment.

(b) Binding cap: if $\sigma^{\text{sep}}(1; \underline{t}) > c$, the upper branch needs to be truncated so that high types pool at $s = c$.

Figure 3: Optimal separating subsidies.

Parameters: $F(x) = x$, $c = 0.5$, $n = 10$. Panel (a): $p = 1.8$. Panel (b): $p = 1$.

Pooling at c creates a trade-off. On the one hand, inspection becomes costless, so any deviation that offers $s < c$ is strictly dominated, even if it claims to be the highest type, and is inspected only after free inspection firms are exhausted. On the other hand, congestion dilutes attention: as more types join the pool, tie-breaking reduce each firm's inspection probability.

Let $q^{\text{pool}}(x) \in (0, 1]$ denote the inspection probability of a firm offering the full subsidy $s = c$ when all types in $[x, 1]$ do so.¹⁵ In equilibrium, the pooling cutoff $\bar{t}$ is determined by two conditions. First, no type $t > \bar{t}$ should prefer to deviate from pooling (**upper IC**). Second, the marginal type $\bar{t}$ must be indifferent between pooling and revealing (**boundary indifference**). These conditions are

$$[t - pc] q^{\text{pool}}(\bar{t}) \geq \sup_{s \in [0, c]} [t - ps] q^{\text{sep}}(\sigma^{-1}(s; \underline{t})), \quad \forall t \in [\bar{t}, 1], \quad (\text{Upper IC})$$

$$[\bar{t} - pc] q^{\text{pool}}(\bar{t}) = [t_1 - p \sigma^{\text{sep}}(\bar{t}; \underline{t})] q^{\text{sep}}(\bar{t}). \quad (\text{Boundary indifference})$$

Here $\sigma^{-1}(\cdot; \underline{t})$ denotes the inverse of the separating schedule $\sigma^{\text{sep}}(\cdot; \underline{t})$.

The **Boundary indifference** condition pin down the marginal type $\bar{t}$. [Lemma 8](#) in [Appendix C.2](#) shows that if $\sigma^{\text{sep}}(1; \underline{t}) > c$, then a unique cutoff $\bar{t}$ always exists. The upper IC constraint follows from the increasing-differences property of the firm's payoff. This cutoff is at the left of the type t_c such that $\sigma^{\text{sep}}(t_c; t_0) = c$ generating a discontinuity in the subsidy policy.

¹⁵Formally,

$$q^{\text{pool}}(x) = \mathbb{E}_K \left[\frac{1 - (1 - \mathbb{E}[t \mid t \geq x])^{K+1}}{(K+1) \mathbb{E}[t \mid t \geq x]} \right],$$

with $K \sim \text{Bin}(n-1, 1 - F(x))$. The derivation is given in [Appendix C.2](#).




Figure 4: Example of a SIS subsidy policy

We now introduce a family of subsidy policies with a step-increasing-step structure. Each member is indexed by the lowest type that is inspected with positive probability.

Definition 5 (Step-increasing-step family). *For any $t_0 \in [\underline{t}, \max\{cp, 1\}]$, define the subsidy policy $\sigma^*(\cdot; t_0) : T \rightarrow [0, c]$ by*

$$\sigma^*(t; \underline{t}) = \begin{cases} 0, & t < t_0, \\ \frac{t}{p} - \frac{\int_{t_0}^t q_n^{\text{sep}}(x; t_0) dx}{p q_n^{\text{sep}}(t; t_0)}, & t \in [t_0, t_1], \\ c, & t > t_1, \end{cases}$$

where $q_n^{\text{sep}}(\cdot; t_0)$ denotes the inspection probability under separation with boundary t_0 , and t_1 is the unique solution to

$$(t_1 - pc) q^{\text{pool}}(t_1) = \int_{t_0}^{t_1} q^{\text{sep}}(x; t_0) dx.$$

Moreover, let $\mu_{t_0}^*$ be a belief system that is consistent with $\sigma^*(\cdot; t_0)$ on path, and off path, satisfies the following: for any subsidy in $(0, t_0/p)$ it assigns belief δ_0 , and for any subsidy in $(\sigma^*(t_1; t_0), c)$ it assigns belief δ_{t_1} . In other words, μ_0 interprets an off-path low subsidy as coming from type 0, and at the discontinuity at the top it assigns belief to the cutoff type t_1 .¹⁶

For $n > 2$, every element of this family exhibits the step-increasing-step pattern: low types offer no subsidy, intermediate types separate, and high types pool at the full subsidy.¹⁷ Moreover, all members of this family are supported on equilibrium.

Lemma 3. *Fix $t_0 \in [\underline{t}, \max\{cp, 1\}]$. The assessment $(\sigma^*(\cdot; t_0), \text{DSIP}_{\mu_{t_0}^*}, \mu_{t_0}^*)$ is an equilibrium.*

¹⁶We write δ_x for the degenerate distribution that places unit mass on type x .

¹⁷When $n = 1$, $q^{\text{sep}}(t; t_0) = \mathbb{1}\{t \geq t_0\}$, so $\sigma^*(t; t_0)$ reduces to a step function: $\sigma^*(t; t_0) = (t_0/p) \mathbb{1}\{t \geq t_0\}$.

Hence the equilibrium set is large and nontrivial.¹⁸ Among this family of equilibria, we single out the the equilibrium $(\sigma^*(\cdot; \underline{t}), DSIP_{\mu_{\underline{t}}^*}, \mu_{\underline{t}}^*)$, which induces inspection of all types down to the lowest possible cutoff $\underline{t}$. We refer to this outcome as the *reasonable equilibrium*. It is efficient because it maximizes the set of types that can be inspected: each inspected type offers a subsidy that is profitable for the firm and simultaneously optimal for the consumer to inspect. As a result, this equilibrium generates the most informative allocation of attention.

4.4 Equilibrium Selection

We now apply a forward-induction refinement, following [Cho and Kreps \(1987\)](#); [Banks and Sobel \(1987\)](#), to narrow the set of equilibria. Off-path beliefs must assign zero probability to types for whom a deviation is dominated. We adopt the equilibrium-dominance restriction. Our main result is that the unique equilibrium outcome surviving this refinement is precisely the one induced by the reasonable equilibrium.

Theorem 2. *For $n \geq 2$, any equilibrium that fails to induce the outcome of $(\sigma^*(\cdot; \underline{t}), DSIP_{\mu_{\underline{t}}^*}, \mu_{\underline{t}}^*)$ fails the equilibrium-dominance refinement.*

To establish the result, consider *optimistic beliefs at s* : the belief system where, after an off-path subsidy s , the posterior assigns probability one to type-1, while every other off-path $s' \neq s$ is assigned probability one to type-0 and never inspected. On-path beliefs remain consistent with σ , so this constitutes an admissible continuation. To highlight the role of s we denote the induced attention by $q_{\sigma, s}$. Under such beliefs, the inspection index reduces to $r_{\mu'}(s) = u - (c - s)$. Hence any $s > c - u$ is inspected with positive probability, although $q_{\sigma, s}(s) < 1$ whenever there exists an on-path s' with $s' > s$ and $(c - s')/(c - s) < \tau_{\mu}(s')$.¹⁹

Lemma 4. *Fix an equilibrium (σ, ι, μ) . Assume there exists an interior pooling at s' , an off-path subsidy s and a type t with $\sigma(t) = s'$ such that, under optimistic beliefs at s , t strictly prefers deviating to s . That is,*

$$(t - ps) q_{\sigma, s}(s) > \pi_{\sigma, \iota}(\sigma(t); t).$$

Then, (σ, ι, μ) fails the equilibrium-dominance refinement.

This lemma provides a simple test for equilibrium selection. We now sketch the proof of uniqueness. The strategy is to rule out, step by step, any equilibrium configuration that fails the equilibrium-dominance refinement under optimistic beliefs. First, we show that any

¹⁸The SIS family does not exhaust the set of equilibria. In particular, there also exist equilibria with interior pooling. In [Appendix D.1](#) we illustrate additional properties that equilibrium subsidy policies may satisfy.

¹⁹If $c > u$, then firms need to offer a high-enough subsidy to induce inspection. Thus, even under optimistic beliefs, some subsidies may never be inspected as the net search cost $c - s$ is still higher than the benefit u .

off-path subsidy s must receive at least as much attention as the lowest on-path subsidy above it would generate under full separation. Using this fact together with the test in [Lemma 4](#), we show that no equilibrium with interior pooling in $[0, c)$ that attracts positive attention can survive. Moreover, a discontinuity can survive only if it occurs at the cap, since any interior jump with attention on both sides would trigger a profitable deviation under optimistic beliefs. Thus the only equilibrium candidates that remain are SIS equilibria. We then establish that the lowest inspected type must coincide with $\underline{t}$, which delivers the uniqueness result.

Fix an equilibrium (σ, ι, μ) . Suppose there exists subsidy levels $s' < s$, where s' is on-path with $q_{\sigma, \iota}(s') > 0$ and s is off-path. Under optimistic beliefs at s , the posterior puts maximal weight on match quality. Hence no lower subsidy, on-path or off-path, can be inspected before s . Since, by assumption, some types choose higher on-path subsidies, s must receive at least the attention given to the lowest such type under full separation. Formally,

$$q_{\sigma, s}(s) \geq \inf_{\{t: \sigma(t) > s\}} q^{\text{sep}}(t; 0).$$

In this worst case, all such higher types also generate a higher inspection index and thus are inspected before.

Now assume that (σ, ι, μ) has an interior pooling level $s \in (0, c)$ that is inspected with positive probability, i.e. $q_{\sigma, \iota}(s) > 0$. Let t^* be the highest type in the pool, with $t^* < 1$. Then σ cannot be right-continuous at t^* .²⁰ In other words, there exists $\varepsilon > 0$ such that every $s' \in (s, s + \varepsilon)$ is off-path. And since $t^* < 1$, there must also exist some higher on-path subsidy. Thus, using the previous logic, under optimistic beliefs for any s' , $q_{\sigma, s'}(s') \geq q^{\text{sep}}(t^*)$. By [Theorem 1](#), inspection probability is increasing in the subsidy. Because there is a mass at s , types just above the pool receive strictly more attention: there exists $\eta > 0$ such that

$$\liminf_{\varepsilon \downarrow 0} [q_{\sigma, \iota}(\sigma(t^* + \varepsilon)) - q_{\sigma, \iota}(s)] \geq \eta.$$

Moreover, the limit from the right can be identified with the separation benchmark. For any type, and in particular for those just above t^* , the worst attention case is that all higher types are inspected with probability one. In that case the inspection probability for such a type coincides with the full-separation benchmark $q^{\text{sep}}(\cdot; 0)$. Since q^{sep} is continuous, taking the limit as $\varepsilon \downarrow 0$ yields

$$\liminf_{\varepsilon \downarrow 0} q_{\sigma, \iota}(\sigma(t^* + \varepsilon)) = q^{\text{sep}}(t^*).$$

²⁰Formally, let $T_\sigma(s) = \{t \in T : \sigma(t) = s\}$, and say that there is pooling at s if $T_\sigma(s)$ has positive F -measure. In any equilibrium with such pooling, σ must be discontinuous at the top of the pool. See [Lemma 14](#) for the general argument.

Therefore, under optimistic beliefs, any off-path s' then secures attention at least as high as $q^{\text{sep}}(t^*)$ which is bounded above from the on-path attention $q_{\sigma,t}(s)$, and thus

$$q_{\sigma,s'}(s') - q_{\sigma,t}(s) \geq \frac{\eta}{2} > 0.$$

Next, note that type t^* must earn positive profit; otherwise, lower types in the pool would prefer to deviate to no subsidy. Hence $t^* - ps > 0$. For any s' close to s , comparison of payoffs yields

$$\begin{aligned} (t^* - ps') q_{\sigma,s'}(s') - (t^* - ps) q_{\sigma,t}(s) &= (t^* - ps) [q_{\sigma,s'}(s') - q_{\sigma,t}(s)] - p(s' - s) q_{\sigma,s'}(s') \\ &\geq (t^* - ps) \frac{\eta}{2} - p(s' - s), \end{aligned}$$

since $0 < q_{\sigma,s'}(s') \leq 1$. Choosing s' sufficiently close to s (for example, $s' = s + (t^* - ps)\eta/(4p)$) makes the payoff strictly higher. Thus t^* has a profitable deviation under optimistic beliefs. We conclude, by [Lemma 4](#), that any equilibrium with a positively inspected pooling below c fails the equilibrium–dominance refinement.

A similar argument rules out any equilibrium with a discontinuity inside an inspected region that does not reach the full subsidy. At such a point, the “jumping” type has a profitable deviation by lowering its subsidy: this both raises its payoff and, under optimistic beliefs, leaves its inspection probability at least unchanged. Since both effects increase profit, the deviation is strictly profitable. The case of a discontinuity at the cap $s = c$ is different. Types pooling at the top are *always* inspected first. If a type in this pool lowers its subsidy, it is inspected only after all top–pool types fail to match, even under optimistic beliefs. In this case, lowering the subsidy strictly reduces inspection probability and may lower profit.

Therefore any equilibrium with an interior pool or an interior discontinuity fails the equilibrium–dominance refinement. The only remaining candidates are the step-increasing-step policies (see [Definition 5](#)). To conclude, we argue that if the lowest inspected type is $t_0 > \underline{t}$, then under optimistic beliefs there is a profitable deviation. Pick $\varepsilon > 0$ so that and $t' := t_0 - \varepsilon > \underline{t}$. Under SIS, t' offers no subsidy, and either it is part of a pooling region with positive inspection, in which our previous argument shows there is a profitable deviation, or it is part of a pool that is never inspected, i.e. $q_{\sigma,t}(0) = 0$. We focus on this second case. Consider the off–path deviation $s' := \underline{t}/p$. Then

$$t' - ps' = t' - \underline{t} > 0,$$

and by construction we also have $q_{\sigma,s'}(s') \geq q^{\text{sep}}(\underline{t}, 0) > 0$. Thus t' has a profitable deviation under optimistic beliefs, so the putative equilibrium fails the equilibrium–dominance refinement.

Together, these lemmas imply that any equilibrium surviving refinement must induce an outcome that coincides with the outcome of the reasonable SIS: no inspection below the participation cutoff, strict revelation on an interval, and a single pooling block at the cap. The bottom cutoff is pinned down by participation at $\underline{t}$, and the top cutoff by boundary indifference at c . By construction, this equilibrium is also the most informative: it has the largest separating region among all equilibria and induces participation from every firm type that can secure positive inspection. Any remaining differences appear only at the boundaries or within the no-inspection region.

5 Welfare and Comparative Statics

We now compute welfare under the reasonable equilibrium, decomposing it into consumer surplus, producer surplus, and total efficiency. We then establish comparative statics with respect to the primitives (p, c, n) .

The next lemma summarizes key properties of the lower and upper cutoff, the attention function, and the separating subsidy policy.

Lemma 5. *The following properties hold:*

- (i) **Lower cutoff.** *The lower cutoff is given by $\underline{t}(p, c) = pc/(1 + p)$. It is strictly increasing in both c and p , and independent of n .*
- (ii) **Attention function.** *For any fixed t_0 , the separating attention function $q_n^{\text{sep}}(t; t_0)$ is strictly decreasing in n for all $t \in (t_0, 1)$. Moreover, for each n and t , $q_n^{\text{sep}}(t; \cdot)$ is decreasing. Therefore, $q_n^{\text{sep}}(\cdot; \underline{t})$ is decreasing in c and p .*
- (iii) **Separating policy.** *For all $t \geq t'_0$, $\sigma_n^{\text{sep}}(t; t'_0) > \sigma_n^{\text{sep}}(t; t_0)$ with $t'_0 > t_0$. In addition, $\sigma_{n+1}^{\text{sep}}(t; t_0) > \sigma_n^{\text{sep}}(t; t_0)$ for all $t > t_0$. We have that $\sigma_n^{\text{sep}}(t; \underline{t})$ is increasing in c and decreasing in p .*
- (iv) **Upper cutoff.** *The upper cutoff $\bar{t}$ is strictly increasing in both p and n .*

A higher inspection cost c or a higher price of subsidies p raises the entry cutoff $\underline{t}$, excluding more low-quality firms. As the number of firms n grows, competition for attention intensifies and the inspection probability of each individual firm falls. To compensate, firms respond with higher subsidies, making the separating schedule steeper. Finally, the upper cutoff $\bar{t}$ increases with p and n , since more expensive subsidies compress the set of types willing to separate before pooling at the cap.

Notation. For convenience, let $q^*(t)$ denote the attention that a type t receives in the reasonable equilibrium: it equals $q^{\text{sep}}(t; \underline{t})$ for $t < \bar{t}$, and $q^{\text{pool}}(\bar{t})$ for $t \geq \bar{t}$.

5.1 Welfare Accounting

Define the average inspection probability per firm as the expected attention a firm generates:

$$Q_n(p, c) := \mathbb{E}_t[q^*(t)] = \int_{\underline{t}}^{\bar{t}} q_n^{\text{sep}}(t) dF(t) + q_n^{\text{pool}}(\bar{t}) [1 - F(\bar{t})]. \quad (9)$$

A type- t firm on the separating branch pays $p \sigma^{\text{sep}}(t; \underline{t})$ upon inspection; envelope arguments yield for $t \in [\underline{t}, \bar{t})$

$$p \sigma^{\text{sep}}(t; \underline{t}) q_n^{\text{sep}}(t; \underline{t}) = t q_n^{\text{sep}}(t; \underline{t}) - \int_{\underline{t}}^t q_n^{\text{sep}}(x; \underline{t}) dx.$$

Pooled types $t \geq \bar{t}$ pay $p c$ upon inspection, so their expected payment equals $p c q_n^{\text{pool}}(\bar{t})$. Thus, on expectation a firm expected transfer costs is given by

$$\begin{aligned} \phi_n^F(p, c) &:= \mathbb{E}_t[p \sigma^*(t) q^*(t)] \\ &= \int_{\underline{t}}^{\bar{t}} \left(t q_n^{\text{sep}}(t; \underline{t}) - \int_{\underline{t}}^t q_n^{\text{sep}}(x; \underline{t}) dx \right) dF(t) + p c q_n^{\text{pool}}(\bar{t}) [1 - F(\bar{t})]. \end{aligned}$$

No match occurs precisely when all inspected firms fail to match which happens with probability $\left(1 - \int_{\underline{t}}^1 x dF(x)\right)^n$.²¹ From this observation, the probability of a match is

$$m_n(p, c) = 1 - \left(1 - \int_{\underline{t}}^1 x dF(x)\right)^n.$$

The *consumer* expected net inspection cost for a single firms equal

$$C_n(p, c) := \mathbb{E}_t[(c - \sigma^*(t)) q^*(t)] = c Q_n(p, c) - \frac{1}{p} \phi_n^F(p, c).$$

Welfare decomposition. Write consumer surplus and producer surplus as

$$CS_n(p, c; t_0) := u m_n(p, c) - n C_n(p, c), \quad (10)$$

$$PS_n(p, c; t_0) := n \mathbb{E}_t[(t - p \sigma^*(t)) q^*(t)] = n \mathbb{E}_t[t q^*(t)] - n \phi_n^F(p, c), \quad (11)$$

²¹The result follows the same logic we use to construct q^{sep} , where now we could think that the outside option is a firm that is inspected after all firms with type above $\underline{t}$ fail to match.

Social welfare adds up consumer and producer surplus:

$$\begin{aligned} W_n(p, c) &:= CS_n(p, c; t_0) + PS_n(p, c; t_0) \\ &= m_n(p, c; t_0) + n(\mathbb{E}_t[(t - c)q^*(t)]) - n\left(1 - \frac{1}{p}\right)\phi_n^F(p, c) \end{aligned} \quad (12)$$

In the special case where $p = 1$ social welfare collapses and results into the expected match value minus the search costs incurred into achieving such match.

5.2 Comparative Statics

We record the main monotonicity properties of the allocation and welfare components.

Theorem 3 (Comparative statics in the reasonable equilibrium). *Then:*

- (i) **Search intensity.** $Q_n(p, c)$ is decreasing in p and in c , and increasing in n .
- (ii) **Match probability.** $m_n(p, c)$ is decreasing in p and in c , and increasing in n .
- (iii) **Consumer surplus.** $CS_n(p, c)$ is decreasing in p and in c , and increasing in n .
- (iv) **Producer surplus.** $PS_n(p, c)$ is decreasing in p and c . Moreover, the per-firm surplus is decreasing in n , i.e. $PS_n(p, c)/n$ decreases in n .

6 Discussion: Agentic Search

The rise of the *agentic economy* is transforming how search operates (Rothschild et al., 2025). Consumers increasingly delegate discovery and evaluation to AI agents that inspect products on their behalf. In this setting, inspection becomes a sequence of auditable computational events (e.g., API calls, token usage, or verification steps) that carry a measurable cost. Attention thus becomes a contractible input that can be priced, transferred, and subsidized.

We embed this development within our framework by modeling a platform that hosts firms and sells *inspection tokens*.²² Each inspection consumes one token at a real cost $c > 0$ to the consumer. Firms can purchase tokens from the platform at price $p \geq 0$ and transfer them to the agent, effectively subsidizing inspection costs. This tokenized approach provides a parsimonious foundation for analyzing efficiency and surplus in emerging agentic marketplaces.

²²A *token* is the basic unit of text processed by a language model, typically a word, part of a word, or punctuation mark, converted into numerical representations the model can manipulate mathematically.

6.1 Platform Revenue Optimization

The platform's revenue problem is setting a token price to maximize revenue. Revenue equals the token price multiplied by the expected token demand, where demand precisely captures the volume of subsidies firms provide. For each price $p \geq 0$, let $(\sigma^*(\cdot; p), \text{DSIP}, \mu^*)$ denote the unique refined outcome of $\mathcal{G}(p)$ under equilibrium-dominance (see [Section 4, Theorem 2](#)). Let $q^*(t; p)$ be the inspection probability induced by DSIP. We define the per-firm *subsidy demand* as:

$$\begin{aligned} D(p) &:= \mathbb{E}[\sigma^*(t; p) q^*(t; p)] \\ &= \int_{\underline{t}(p)}^{\bar{t}(p)} \sigma^{\text{sep}}(t) q^{\text{sep}}(t; \underline{t}(p)) dF(t) + c q^{\text{pool}}(\bar{t}(p)) [1 - F(\bar{t}(p))], \end{aligned}$$

where $\underline{t}(p) = \frac{pc}{1+p}$ is the participation cutoff and $\bar{t}(p)$ the pooling cutoff. Platform revenue is then:

$$R(p) = n p D(p).$$

Let $\psi(t) = t - \frac{1-F(t)}{f(t)}$ be the virtual value function. Per-firm revenue admits the decomposition:

$$\frac{R(p)}{n} = \underbrace{\left(\int_{\underline{t}(p)}^{\bar{t}(p)} \psi(x) q^{\text{sep}}(x; \underline{t}(p)) \frac{dF(x)}{F(\bar{t}(p))} \right)}_{\text{separating branch}} F(\bar{t}(p)) + \underbrace{\left(\bar{t}(p) q^{\text{pool}}(\bar{t}(p)) \right)}_{\text{pooling branch}} [1 - F(\bar{t}(p))]. \quad (13)$$

This decomposition reveals two distinct revenue components.²³ The first term represents the *revealing branch*, where firms separate by type and revenue corresponds to expected virtual surplus weighted by inspection probabilities. This component exhibits a single-peaked relationship with price: lower prices expand subsidies and increase inspection, but beyond the unit-elasticity point, further price cuts reduce revenue. At high prices, the pooling region vanishes and only revelation contributes to revenue.

The second term captures the *pooling branch*, where all types above $\bar{t}(p)$ are grouped and effectively pay the cutoff virtual value $\bar{t}(p)$. This component increases monotonically in price: higher prices raise per-unit revenue from the pooled types, though the pool's relative weight shrinks as participation decreases.

Crucially, price does not merely scale revenue, it selects the *margin of competition*. Lower prices reduce per-unit revenue but broaden participation, and increase subsidies, further it expands the pooling region, giving greater weight to the pooling branch. Higher prices

²³See [Lemma 17](#) in the appendix to see how we reach such representation.

shrink participation and thin the pool, though each inspected type provides larger subsidies. In this regime, the pooling branch matters less and revenue is primarily driven by the separating branch. Overall, revenue combines a hump-shaped revealing branch with a monotone pooling branch. The unique revenue-maximizing price p^* emerges where these opposing forces balance.

Theorem 4 (Excess Search). *Suppose F satisfies the regularity condition. Then the platform's revenue maximization problem admits a unique solution $p^* > 0$. At this optimal price:*

1. **Pooling remains active:** $\bar{t}(p^*) < 1$.
2. **Negative virtual values receive inspection:** *Some types with $\psi(t) < 0$ are inspected with positive probability.*

The platform optimally sets p^* to thicken the top pool, steering attention toward inspections that a virtual surplus benchmark would exclude. Since generating additional attention is cheapest in this pooling region, profit maximization occurs precisely when the market is induced to search excessively. This creates a fundamental tension in agentic marketplaces. While consumers benefit from increased match likelihood and reduced net inspection costs, attention becomes distorted toward excessive inspection, obfuscating discovery. The platform's revenue maximization thus generates a welfare loss: the market searches too much relative to the first-best allocation.

7 Concluding Remarks

This paper develops a theory of attention markets in which firms can directly subsidize costly search. In such environments, subsidies play a dual role. They serve as signals, conveying information about a firm's expected match value. At the same time, they are instrumental, reducing the consumer's effective inspection costs and thereby altering the sequence of search. In equilibrium, these two roles jointly determine which firms participate, which reveal their type, and which pool. In this way, subsidies endogenously structure the allocation of attention.

We show that every equilibrium satisfies a *subsidy-sorting principle*: higher types weakly offer higher subsidies, and the consumer follows a descending-subsidy inspection order. Within this structure, the unique equilibrium outcome surviving refinement is the *reasonable equilibrium*. It takes a step-increasing-step form: no inspection below a participation cutoff, revelation on an intermediate region, and a pooling block at the cap. This equilibrium is efficient in the sense of being most informative: it maximizes the set of types that can induce

inspection, subject to incentive constraints. Any information that is lost is either uninspected altogether or free to obtain.

These results provide a benchmark for platform design. If a platform can mint and sell inspection tokens, its problem reduces to a linear-cost mechanism. Yet profit motives may distort attention. By thickening the top pool, a platform can steer consumers to inspect too much, raising revenue but lowering social efficiency.

The framework also delivers empirical predictions. Subsidies should be weakly increasing in observable proxies for match quality. High-subsidy pools should exhibit congestion, with inspection probabilities falling as the pool thickens. Changes in token prices should shift the participation margin and alter match rates and transfers in predictable ways. Such patterns can be measured in platform data, offering a path for empirical validation.

More broadly, the analysis highlights how inspection can be organized without auctions. Direct subsidies decentralize the allocation of attention, sustaining efficiency so long as ranking rules remain transparent. As AI agents increasingly intermediate search, this principle provides both a microfoundation for programmable attention and a policy guide for preserving efficiency in environments where the temptation to steer is strongest.

References

- Anderson, Simon P. and Régis Renault**, “Advertising Content,” *American Economic Review*, March 2006, 96 (1), 93–113.
- Athey, Susan and Glenn Ellison**, “Position Auctions with Consumer Search,” *The Quarterly Journal of Economics*, 2011, 126 (3), 1213–1270.
- Au, Pak Hung and Mark Whitmeyer**, “Attraction versus Persuasion: Information Provision in Search Markets,” *Journal of Political Economy*, 2023, 131 (1), 202–245.
- Banks, Jeffrey S and Joel Sobel**, “Equilibrium selection in signaling games,” *Econometrica*, 1987, 55 (3), 647–661.
- Board, Simon and Jay Lu**, “Competitive Information Disclosure in Search Markets,” *Journal of Political Economy*, 2018, 126 (5), pp. 1965–2010.
- Chade, Hector and Lones Smith**, “Simultaneous search,” *Econometrica*, 2006, 74 (5), 1293–1307.
- Chen, Yongmin and Chuan He**, “Paid Placement: Advertising and Search on the Internet,” *The Economic Journal*, 2011, 121 (556), F309–F328.
- Cho, In-Koo and David M Kreps**, “Signaling Games and Stable Equilibria,” *The Quarterly Journal of Economics*, 1987, 102 (2), 179–221.
- Diamond, Peter A**, “A model of price adjustment,” *Journal of Economic Theory*, 1971, 3, 156–168.
- Ding, Bolin, Yiding Feng, Chien-Ju Ho, Wei Tang, and Haifeng Xu**, “Competitive Information Design for Pandora’s Box,” 2025.
- Edelman, Benjamin, Michael Ostrovsky, and Michael Schwarz**, “Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords,” *American Economic Review*, 2007, 97 (1), 242–259.

- Ellison, Glenn and Alexander Wolitzky**, “A search cost model of obfuscation,” *The RAND Journal of Economics*, 2012, 43 (3), 417–441.
- Janssen, Maarten, Thomas Jungbauer, Marcel Preuss, and Cole Williams**, “Search Platforms: Big Data and Sponsored Positions,” 2025.
- Kreps, David M and Robert Wilson**, “Sequential equilibria,” *Econometrica*, 1982, pp. 863–894.
- McCall, J J**, “Economics of Information and Job Search,” *The Quarterly Journal of Economics*, 1970, 84 (1), 113–126.
- Mekonnen, Teddy, Zeky Murra-Anton, and Bobak Pakzad-Hurson**, “Persuaded Search,” *Journal of Political Economy*, 2025, 133 (10), 3167–3207.
- Milgrom, Paul and Chris Shannon**, “Monotone Comparative Statics,” *Econometrica*, 1994, 62 (1), 157–180.
- Nocke, Volker and Patrick Rey**, “Consumer Search, Steering, and Choice Overload,” *Journal of Political Economy*, May 2024, 132 (5), 1684–1739.
- Rochet, Jean-Charles and Jean Tirole**, “Platform Competition in Two-Sided Markets,” *Journal of the European Economic Association*, 2003, 1 (4), 990–1029.
- Rothschild, David M., Markus Mobius, Jake M. Hofman, Eleanor W. Dillon, Daniel G. Goldstein, Nicole Immorlica, Sonia Jaffe, Brendan Lucier, Aleksandrs Slivkins, and Matthew Vogel**, “The Agentic Economy,” 2025.
- Stigler, George J.**, “The Economics of Information,” *Journal of Political Economy*, 1961, 69 (3), 213–225.
- Weitzman, Martin L.**, “Optimal Search for the Best Alternative,” *Econometrica*, 1979, 47 (3), 641–654.
- Weyl, E Glen**, “A price theory of multi-sided platforms,” *American Economic Review*, 2010, 100 (4), 1642–1672.

A A price competition Model

In this appendix, we consider an alternative version of the model where firms do not subsidize search, rather they announce prices and then the consumer use those to determine how to search and then inspect.

We show that as in the baseline model, there are multiple equilibria and they satisfy either pooling or a equal revenue condition. Moreover, we show that our equilibrium refinement do not have a bite in this environment. That is, a forward induction argument do not single out any equilibrium outcome. As we argue, this follows as this version of the game lacks a costly signaling structure.

A.1 Framework

Consider a set of firms and one consumer. All firms offer a product, and the consumer values the product of firm j by $v_j \in \{0, v\}$ where the firm knows that $\Pr\{v_j = v\} = t_j$. However, the consumer only knows that $v_j \sim F$ i.i.d.

The firms all observe their private info and determine a price for their product, denoted by $p_j \in [0, v]$. The consumer observe the prices, and then determine a search sequence. The consumer need to costly inspect each product to realize the value she derives from it.

A.2 Multiple pooling equilibrium

First, we show that there exists a lot of pooling equilibria. Assume that all firms pool and charge a price $p > 0$. Then, the consumer needs to inspect products randomly. Moreover, assume that if the consumer observes a price $p' \neq p$, she believes it was the lowest type who propose it. We observe that if

$$(v - p)\mathbb{E}[t] > c,$$

then the consumer will search randomly until either inspecting all firms or finding the first firm with value $v_j = v$. She then buys the object at price p and stops searching.

For firms, if they follow the strategy, they guarantee some positive expected payoff. However, if the firm decides to deviate and offer a different price, the consumer will never inspect such item (as it is just expected to be a cost for her), and thus the firm will generate a payoff of zero.

No one wants to deviate, and thus. In this case, we have infinite outcome of equilibrium outcomes.

B Omitted proofs for Section 3

B.1 Proof of Lemma 1

First, we establish that the attention function is weakly increasing on path.

Claim 1. For $s, s' \in \text{Im}(\sigma)$, if $s' \geq s$, then $q_{\sigma, \iota}(s') \geq q_{\sigma, \iota}(s)$.

Proof. Assume towards a contradiction that this is not the case, then $q_{\sigma, \iota}(s') < q_{\sigma, \iota}(s)$. Let t' be such that $\sigma(t') = s'$, then from the firm optimality it should hold that

$$(t' - ps')q_{\sigma, \iota}(s') \geq (t - ps)q_{\sigma, \iota}(s)$$

and in particular, for s . However, note that

$$(t' - ps')q_{\sigma, \iota}(s') < (t - p\sigma(t))q_{\sigma, \iota}(s) < (t' - ps)q_{\sigma, \iota}(s).$$

Thus, type t' has a profitable deviation. □

Note the proof establish something stronger. In particular, we have that if $s \in \text{Im}(\sigma)$, then for all $s' < s$, $q_{\sigma, \iota}(s') \leq q_{\sigma, \iota}(s)$ regardless of s' being on-path or off-path.

Second, we show that the firm payoff over inspection probabilities satisfies single-crossing differences in type. Thus, higher types firms are willing to pay more for the same level of attention. In turn, this means that the subsidy policy should be increasing.

Claim 2. The subsidy policy is increasing: if $t' \geq t$, then $\sigma(t') \geq \sigma(t)$.

Proof. Let $q(s)$ denote the ex-ante probability that a firm offering subsidy s is inspected in equilibrium. Define $q(t) := q(\sigma(t))$ for shorthand. Since $\sigma(t)$ is an equilibrium strategy, it must solve:

$$\sigma(t) \in \arg \max_{s \in [0, c]} [t - ps] \cdot q(s),$$

which implies $[t - p\sigma(t)] \geq 0$ since $s = 0$ is always feasible and $t, q(\cdot) \in [0, 1]$

For any two types $t' > t$, it holds as follows:

$$[t - p\sigma(t)] \cdot q(t) \geq [t - p\sigma(t')] \cdot q(t'), \tag{14}$$

$$[t' - p\sigma(t')] \cdot q(t') \geq [t' - p\sigma(t)] \cdot q(t). \tag{15}$$

Re-arranging (14) and (15), we obtain:

$$\begin{aligned} p \cdot [\sigma(t')q(t') - \sigma(t)q(t)] &\geq t \cdot [q(t') - q(t)], \\ t' \cdot [q(t') - q(t)] &\geq p \cdot [\sigma(t')q(t') - \sigma(t)q(t)]. \end{aligned}$$

Combining the two inequalities implies

$$t'[q(t') - q(t)] \geq t[q(t') - q(t)],$$

or equivalently,

$$(t' - t) \cdot [q(t') - q(t)] \geq 0.$$

Because $t' > t$, it follows that $q(t') \geq q(t)$. That is, inspection probabilities are weakly increasing in type.

Finally, since the inspection function $q(\cdot)$ is weakly increasing in subsidies (by the previous claim), it must hold that $\sigma(t') \geq \sigma(t)$. In other words, equilibrium subsidies are monotone in type. $\square$

B.2 Proof of Lemma 2

Proof. By Weitzman's reservation-value theorem, the consumer's optimal inspection policy is given by the index rule: inspect firms in decreasing order of their reservation (index) values. If several firms share the same index, any symmetric tie-breaking rule yields the same expected payoff; under the symmetry assumption the only way to break ties is to randomize uniformly among tied firms. Hence the index rule with uniform tie-breaking (when needed) is optimal. $\square$

B.3 Proof of Corollary 1

Proof. On-path higher subsidies correspond to higher types, moreover higher subsidies also decrease the search cost. Both forces imply that search a higher subsidy is associated with a higher reservation index. Thus, on-path the consumer searches in descending order of subsidy. The optimal stopping point is determined when $r_\mu(s) < 0$. That is,

$$u - \frac{c - s}{\tau_\mu(s)} < 0 \iff u\mathbb{E}[t \mid \sigma(t) = s] < c - s$$

$\square$

B.4 Proof of Theorem 1

Proof. The result immediately follows from Lemma 1 and Lemma 2. In its current statement, we used the result in Corollary 1. $\square$

C Technical details and supplementary analysis for [Section 4](#)

C.1 Inspection probability under full separation

Lemma 6 (Inspection probability under full separation). *Suppose the subsidy policy σ is strictly increasing and the consumer follows the descending-subsidy index rule $DSIP_\mu$ where μ is on-path consistent with σ . Let $\underline{t} \in [0, c)$ denote the lowest inspected type, defined by $\underline{t} = \inf\{t : \sigma(t) = c - ut\}$. Then, for any $t \in T$, the ex-ante inspection probability of a type- t firm is equivalent to*

$$q_n^{sep}(t; \underline{t}) = \left(1 - \int_t^1 x dF(x)\right)^{n-1} \mathbf{1}\{t \geq \underline{t}\}.$$

The attention function $q_n^{sep}(\cdot; \underline{t})$ is strictly increasing on $(\underline{t}, 1)$, hence differentiable a.e., and $q_n^{sep}(1; \underline{t}) = 1$.

Proof. First, under σ strictly increasing, and the inspection policy following $DSIP_\mu$, the consumer inspect all firms that offer a subsidy level $s \geq \sigma(\underline{t})$.

For $x \geq \underline{t}$, consider a type x -firm. If such firm offers the k -th highest subsidy (rank k), its probability of inspection is given by expected failure rate $\prod_{\ell=1}^{k-1} (1 - t_{(\ell:n-1)})$ when $t_{(\ell:n-1)} \geq x$, where $t_{(\ell:n-1)}$ is the ℓ -th highest type among the other $n-1$ firms. Hence the ex-ante inspection probability is given by

$$q_{\sigma, DSIP_\mu}(x) = \sum_{k=1}^n \mathbb{E} \left[\prod_{\ell=1}^{k-1} (1 - t_{(\ell:n-1)}) \mid t_{(k-1:n-1)} \geq x > t_{(k:n-1)} \right] \Pr\{t_{(k-1:n-1)} \geq x > t_{(k:n-1)}\},$$

with the convention $t_{(0:m)} := 1$, and $q_{\sigma, DSIP_\mu}(x) = 0$ for $x < \underline{t}$.

Fix x . For $k = 1, \dots, N$, consider the event

$$\{t_{(k-1:n-1)} \geq x > t_{(k:n-1)}\},$$

which is exactly the event that $k-1$ out of the $N-1$ opponent types lie above x . Since each opponent exceeds x with probability $1 - F(x)$, the count K of such types follows a binomial law: $K \sim \text{Bin}(N-1, 1 - F(x))$. Thus

$$\Pr\{t_{(k-1:n-1)} \geq x > t_{(k:n-1)}\} = \Pr\{K = k-1\} = \binom{N-1}{k-1} (1 - F(x))^{k-1} F(x)^{N-k}.$$

Conditional on this event, the $k-1$ opponents above x are i.i.d. draws from the truncated

distribution $F(\cdot \mid t \geq x)$. Hence,

$$\mathbb{E} \left[\prod_{\ell=1}^{k-1} (1 - t_{(\ell:n-1)}) \mid t_{(k-1:n-1)} \geq x > t_{(k:n-1)} \right] = \left(\mathbb{E}[1 - t \mid t \geq x] \right)^{k-1}.$$

Therefore,

$$\begin{aligned} q_{\sigma, DSIP_{\sigma}}(x) &= \sum_{k=1}^n (\mathbb{E}[1 - t \mid t \geq x])^{k-1} \binom{n-1}{k-1} (1 - F(x))^{k-1} F(x)^{n-k} \\ &= \sum_{k=0}^{n-1} \binom{n-1}{k} \left((1 - F(x)) (\mathbb{E}[1 - t \mid t \geq x]) \right)^k F(x)^{n-1-k}. \end{aligned}$$

The binomial theorem now gives

$$q_{\sigma, DSIP_{\mu}}(x) = \left[F(x) + (1 - F(x)) - \int_x^1 t dF(t) \right]^{n-1} = \left(1 - \int_x^1 t dF(t) \right)^{n-1} = q_n^{\text{sep}}(x; \underline{t}),$$

where we use $\mathbb{E}[1 - t \mid t \geq x] = 1 - \frac{1}{1-F(x)} \int_x^1 t dF(t)$. Establishing the result. The remaining properties follow directly from inspection. $\square$

C.2 Inspection probability under pooling at the top

Lemma 7. *If all types in $[t_1, 1]$ pool at c , then the ex-ante inspection probability for a pooled firm is*

$$q_n^{\text{pool}}(t_1) = \mathbb{E}_K \left[\frac{1 - (1 - \mathbb{E}[t \mid t \geq t_1])^{K+1}}{(K+1) \mathbb{E}[t \mid t \geq t_1]} \right],$$

where $K \sim \text{Bin}(n-1, 1 - F(t_1))$. Moreover, q^{pool} is continuous and strictly increasing.

Proof. Let K be the number of *other* firms that also lie in $[t_1, 1]$; by i.i.d. types, $K \sim \text{Bin}(n-1, 1 - F(t_1))$. Conditional on $K = k$, the pool size is $m = k + 1$ (including the focal firm). Uniform tie-breaking makes the inspection order among the m pooled firms uniformly random. Fix the focal firm and suppose it is the j -th in that random order. The consumer reaches it if, and only if, all $j - 1$ predecessors fail to match. Conditional on being in the pool, types are i.i.d. with the conditional law $F(\cdot \mid t \geq t_1)$; hence by independence,

$$\Pr\{\text{all } j - 1 \text{ predecessors fail} \mid t \geq t_1\} = \left(\mathbb{E}[1 - t \mid t \geq t_1] \right)^{j-1} = (1 - \tau(t_1))^{j-1}.$$

Averaging over the uniformly random position $j \in \{1, \dots, m\}$,

$$\Pr\{\text{focal is inspected} \mid K = k\} = \frac{1}{m} \sum_{j=1}^m (1 - \tau(t_1))^{j-1} = \frac{1 - (1 - \tau(t_1))^m}{m \tau(t_1)}.$$

Finally, take expectations over K to obtain the stated $q^{\text{pool}}(t_1)$. Continuity follows from F having full support, and strict increasingness from $\mathbb{E}[t \mid t \geq t_1]$ being increasing in t_1 . $\square$

C.3 Omitted proofs for [Section 4](#)

Lemma 8. Fix $t_0 \in T_0(p, c)$. Define the upper cutoff $t_1 : T_0(p, c) \rightarrow [pc, 1]$ as follows: if $\sigma^*(1; t_0) < c$, set $t_1 = 1$; otherwise, let

$$H(x; t_0) = q^{\text{pool}}(x)(x - pc) - q^{\text{sep}}(x; t_0)(x - p\sigma^{\text{sep}}(x; t_0)),$$

and define t_1 as the unique solution to $H(t_1; t_0) = 0$. The mapping $t_0 \mapsto t_1(t_0)$ is well defined, continuous, and differentiable on $T_0(p, c)$.

Proof. We begin with boundary values of $H(\cdot; t_0)$. At $x = pc$,

$$H(pc; t_0) = -q^{\text{sep}}(pc; t_0)(pc - p\sigma^{\text{sep}}(pc; t_0)) = -\int_{t_0}^{pc} q^{\text{sep}}(y; t_0) dy < 0.$$

At $x = t_c$ (with $\sigma^{\text{sep}}(t_c; t_0) = c$),

$$H(t_c; t_0) = (t_c - pc)(q^{\text{pool}}(t_c) - q^{\text{sep}}(t_c; t_0)) > 0,$$

as for all x , $q^{\text{pool}}(x) > q^{\text{sep}}(x; t_0)$ (which also implies that $H(x; t_0) > 0$ for all $x > t_c$). Since $H(\cdot; t_0)$ is continuous, the intermediate value theorem guarantees at least one root $t_1 \in (pc, t_c)$.

Uniqueness. To establish uniqueness, we show that $H(\cdot; t_0)$ has a single crossing. Recall the envelope identity

$$q^{\text{sep}}(x; t_0)(x - p\sigma^{\text{sep}}(x; t_0)) = \int_{t_0}^x q^{\text{sep}}(y; t_0) dy.$$

It follows that

$$H(x; t_0) > 0 \iff q^{\text{pool}}(x) > \frac{\int_{t_0}^x q^{\text{sep}}(y; t_0) dy}{x - pc} =: H_-(x).$$

Differentiating yields

$$H'_-(x) = \frac{p q^{\text{sep}}(x; t_0)}{(x - pc)^2} (\sigma^{\text{sep}}(x; t_0) - c).$$

Then $H'_-(t_c) = 0$, $H'_-(x) < 0$ for $x < t_c$, and $H'_-(x) > 0$ for $x > t_c$. Since $q^{\text{pool}}(x)$ is strictly increasing, then $q^{\text{pool}}(x) - H'_-(x)$ is strictly decreasing. Since we already establish the existence of a zero, this guarantees uniqueness.

Hence $H(\cdot; t_0)$ has a unique root t_1 , which establishes existence and uniqueness of the cutoff. Moreover, such root is to the left of t_c . The implicit function theorem establishes that that $t_1(t_0)$ is differentiable, where

$$t'_1(t_0) = -\frac{q^{\text{sep}}(t_0; t_0)(t_1 - pc)}{(q^{\text{pool}})'(t_1)(t_1 - pc)^2 + pq^{\text{sep}}(t_1)(c - \sigma^{\text{sep}}(t_1; t_0))} < 0,$$

where we use that $pc < t_1 < t_c$, and thus $c > \sigma^{\text{sep}}(t_1; t_0)$. $\square$

Proof of Lemma 3. Fix $t_0 \in T_0(p, c)$ and consider the assessment $(\sigma^*(\cdot; t_0), \text{DSIP}_{\mu_{t_0}^*}, \mu_{t_0}^*)$. By Lemma 2, for any belief system μ the descending-subsidy index policy DSIP_μ maximizes the consumer's expected utility. By construction, $\mu_{t_0}^*$ coincides with Bayes' rule on the realized support of $\sigma^*(\cdot; t_0)$. Hence the consumer's strategy is optimal and beliefs are on-path consistent. Types $t \leq t_0$ pool at $s = 0$. Their equilibrium payoff is

$$\pi^*(0; t) = (t - p \cdot 0) q_{\sigma^*, \text{DSIP}}(0) \geq 0,$$

since $q_{\sigma^*, \text{DSIP}}(0) \geq 0$. Any deviation that raises attention must jump to at least the bottom of the revealing branch, $s \geq \sigma^*(t_0; t_0) = \frac{t_0}{p}$, because lower subsidies do not improve rank under $\text{DSIP}_{\mu_{t_0}^*}$. For $t \leq t_0$ this yields

$$\pi(s; t) = (t - ps) q_{\sigma^*, \text{DSIP}}(s) \leq (t - t_0) q_{\sigma^*, \text{DSIP}}(s) \leq 0,$$

with strict inequality if $s > \frac{t_0}{p}$. Thus no $t \leq t_0$ gains by deviating. For $t \in (t_0, \bar{t})$ the policy $\sigma^*(\cdot; t_0)$ is strictly increasing and solves the standard envelope (local IC) condition in (7), so truth telling maximizes interim profit along the branch. A deviation to a lower subsidy reduces rank and attention under $\text{DSIP}_{\mu_{t_0}^*}$; a deviation to a higher subsidy cannot improve attention at that rank but strictly raises the per-inspection payment ps . Either way the deviation weakly lowers profit, with strict loss whenever rank changes or s increases.

Combining, the consumer is best responding given $\mu_{t_0}^*$, beliefs are on-path Bayes consistent, and no firm type has a profitable deviation. Hence $(\sigma^*(\cdot; t_0), \text{DSIP}_{\mu_{t_0}^*}, \mu_{t_0}^*)$ is a symmetric PBE. $\square$

D Equilibrium Properties and the Refinement

This appendix collects general properties of equilibria and clarifies how the equilibrium-dominance refinement operates. These results underpin the arguments in [Section 4.4](#).

By [Theorem 1](#), in any equilibrium (σ, ι, μ) the subsidy policy σ is weakly increasing and the consumer follows DSIP $_{\mu}$ ([Definition 4](#)). Let $q_{\sigma}(s)$ denote the induced attention function for subsidy s .

Notation. Let $S = [0, c]$ be the set of feasible subsidies, $S_{\sigma} = \text{Im}(\sigma)$ the on-path set, and $S_{\sigma}^{-} = S \setminus S_{\sigma}$ the off-path set. For $s \in S_{\sigma}$, let $T_{\sigma}(s) = \{t : \sigma(t) = s\}$ be the set of types mapped to s . Let $q^{\text{sep}}(t) = (1 - \int_t^1 x dF(x))^{n-1}$ be the inspection probability under a strictly separating branch (derived in [Appendix C.1](#)), and let $q^{\text{pool}}(\bar{t})$ be the inspection probability of a pooled firm when all types in $[\bar{t}, 1]$ pool at c (derived in [Appendix C.2](#)).

D.1 General Properties of Equilibria

We first establish several properties that hold in any equilibrium.

- (i) In any inspected region, at most one type earns zero profit.
- (ii) If there is interior pooling, the subsidy policy fails right-continuity.
- (iii) On a strictly separating branch, types must satisfy the differential equation derived earlier, with the appropriate initial condition.

Among types that receive attention, at most one can break even. Higher types must earn strictly positive rents.

Lemma 9. *Fix an equilibrium (σ, ι, μ) . If there exist types $t_1 < t_2$ with $q_{\sigma}(\sigma(t_1)), q_{\sigma}(\sigma(t_2)) > 0$, then $\pi_{\sigma, \iota}(\sigma(t_2); t_2) > 0$.*

Proof. Fix two types $t_1 < t_2$ that both receive positive attention. By [Theorem 1](#), $q(\cdot)$ is weakly increasing in S_{σ} , and firms face payoffs $\pi(s; t) = (t - ps)q(s)$. If $\pi(\sigma(t_1), t_1) = (t - p\sigma(t_1))q_{\sigma}(\sigma(t_1)) = 0$, then $t_1 = p\sigma(t_1)$. Since $t_2 > t_1$, it follows that $\pi(\sigma(t_1); t_2) > 0$, and by optimality $\pi(\sigma(t_2), t_2) \geq \pi(\sigma(t_1); t_2) > 0$. Thus at most one inspected type breaks even; all others either earn rents or are not inspected. $\square$

The next lemma shows that a pool inside $(0, c)$ forces a discontinuity: the attention probability jumps upward just to the right of s , so a highest pooled type would deviate.

Lemma 10. *Fix an equilibrium (σ, ι, μ) . If there exists $s \in (0, c)$ such that $\Pr\{t \in T_{\sigma}(s)\} > 0$, then there exists $\varepsilon > 0$ with $(s, s + \varepsilon) \subseteq S_{\sigma}^{-}$.*

Proof. Suppose, toward a contradiction, that $[s, s + \varepsilon) \subseteq S_\sigma$. We show q_σ has a strict right-hand jump at s .

Without loss of generality assume σ^{-1} is strictly increasing on $(s, s + \varepsilon)$ (otherwise restrict the interval). For any $s' \in (s, s + \varepsilon)$, a firm at s' is ranked strictly above those at s and is inspected with probability $q_\sigma(s') = q^{\text{sep}}(\sigma^{-1}(s'))$. A firm at s faces pooling. Let K be the number of other pooled firms at s , with $\Pr\{K \geq 1\} > 0$ since $\Pr\{t \in T_\sigma(s)\} > 0$. Conditional on reaching level s and on $K = k \geq 1$, its inspection probability is

$$q_{\text{pool}}(k) = \frac{1 - (1 - \tau)^{k+1}}{(k+1)\tau} < 1,$$

where τ is the common pooled posterior success probability. Hence

$$q_\sigma(s) = q^{\text{sep}}(\sigma^{-1}(s)) \mathbb{E}[q_{\text{pool}}(K)] < q^{\text{sep}}(\sigma^{-1}(s)).$$

Taking $s' \downarrow s$ and using continuity of $q^{\text{sep}} \circ \sigma^{-1}$ on-path,

$$\lim_{s' \downarrow s} (q_\sigma(s') - q_\sigma(s)) = q^{\text{sep}}(\sigma^{-1}(s)) (1 - \mathbb{E}[q_{\text{pool}}(K)]) > 0,$$

since $\mathbb{E}[q_{\text{pool}}(K)] < 1$ and $q^{\text{sep}}(\sigma^{-1}(s)) > 0$.

Let $t^* = \sup T_\sigma(s)$; if $t^* \notin T_\sigma(s)$ pick a sequence (t^k) in $T_\sigma(s)$ approaching t^* . Then for s' close to s ,

$$\begin{aligned} \pi(t^*; s') - \pi(t^*; s) &= (t^* - ps')q_\sigma(s') - (t^* - ps)q_\sigma(s) \\ &\geq (t^* - ps)(q_\sigma(s') - q_\sigma(s)) - p(s' - s), \end{aligned}$$

Hence, as

$$\lim_{s' \rightarrow s} (\pi(t^*; s') - \pi(t^*; s)) = (t^* - ps)q^{\text{sep}}(\sigma^{-1}(s)) (1 - \mathbb{E}[q_{\text{pool}}(K)]) > 0$$

a contradiction. Hence $(s, s + \varepsilon) \subseteq S_\sigma^-$ for some $\varepsilon > 0$. $\square$

Finally, we show that on any separating branch the ODE at (7) precisely pins down the subsidy policy.

Lemma 11. *Let (σ, ι, μ) be an equilibrium. Suppose there exists an interval (t_1, t_2) on which σ is strictly increasing and $q_\sigma(\sigma(t_1)) > 0$. Then on (t_1, t_2) σ solves (7) on (t_1, t_2) , that is*

$$\sigma'(t) = \frac{(t - p\sigma(t)) q^{\text{sep}'}(t)}{p q^{\text{sep}}(t)}.$$

Proof. Fix $t \in (t_1, t_2)$. Since σ is strictly increasing on (t_1, t_2) , subsidies identify types on this interval, and the induced attention along the branch satisfies for $x \in (t_1, t_2)$, $q_\sigma(\sigma(x)) = q^{\text{sep}}(x)$. At an interior optimum $s = \sigma(t)$ the first-order condition holds:

$$0 = \frac{d}{ds} \pi(s; t) \Big|_{s=\sigma(t)} = -p q_\sigma(\sigma(t)) + (t - p \sigma(t)) \frac{d}{ds} q_\sigma(s) \Big|_{s=\sigma(t)}.$$

Along the revealing branch $q_\sigma(s) = q^{\text{sep}}(\sigma^{-1}(s))$, so by the chain rule,

$$\frac{d}{ds} q_\sigma(s) \Big|_{s=\sigma(t)} = \frac{q^{\text{sep}'}(t)}{\sigma'(t)}.$$

Substituting into the FOC, using $q_\sigma(\sigma(t)) = q^{\text{sep}}(t)$ and rearranging yields

$$\sigma'(t) = \frac{(t - p \sigma(t)) q^{\text{sep}'}(t)}{p q^{\text{sep}}(t)}.$$

Because $q_\sigma(\sigma(t_1)) > 0$, the optimum is interior in a right neighborhood of t_1 , so the FOC applies there; strict monotonicity extends it throughout (t_1, t_2) . This coincides with (7). $\square$

D.2 Equilibrium-Dominance and its implications

Equilibrium-Dominance definition. For an off-path s , a continuation (ι', μ') is *admissible* at s , $(\iota', \mu') \in \mathcal{C}(s)$, if (i) ι' is DSIP $_{\mu'}$, and (ii) on-path beliefs are Bayes-consistent: for $s \in S_\sigma$, $\mu(t \mid s) = \frac{f(t)}{\int_{T_\sigma(s)} f(x) dx}$. A type- t never wants to deviate and offer s , $t \in D_{\sigma, \iota}(s)$ if for any admissible continuation ,

$$(t - ps) q_{\sigma, \iota'}(s) < (t - p\sigma(t)) q_{\sigma, \iota'}(\sigma(t)) = \pi_{\sigma, \iota}(\sigma(t); t),$$

An equilibrium (σ, ι, μ) *fails equilibrium-dominance* if there exists $s \in S_\sigma^-$ and $t \in D_{\sigma, \iota}(s)$ with $t \in \text{supp}(\mu(s))$.

Optimistic beliefs. For some off-path s , we refer as *optimistic beliefs at s* the belief system where subsidy s is assigned a degenerate “type-1” posterior, while every other off-path $s' \neq s$ is assigned a degenerate “type-0” posterior and never inspected. On-path beliefs remain consistent with σ , so this constitutes an admissible continuation. Under such beliefs, the inspection index reduces to $r_{\mu'}(s) = u - (c - s)$. Hence any $s > c - u$ is inspected with positive probability, though $q_{\sigma, \iota'}(s) < 1$ whenever there exists $s' \in S_\sigma$ with $s' > s$ and $(c - s')/(c - s) < \tau_{\mu'}(s')$.

Proof of Lemma 4. Optimistic beliefs at s are an admissible continuation by construction. Moreover, this belief system places the higher match posterior at s , thus yields the highest possible

attention that is consistent on path with σ . That is, for any $(t', \mu') \in \mathcal{C}(s)$, $q_{\sigma, s}(s) \geq q_{\sigma, \mu'}(s)$. Thus if

$$(t - ps)q_{\sigma, s}(s) \geq (t - p\sigma(t))q_{\sigma, \iota}(\sigma(t))$$

Any type $t' > t$ such that $\sigma(t) = s' = \sigma(t')$ benefits more from such deviation as π satisfies increasing differences in s and q . Thus, it should hold that $\mu(s) = \sup\{t : \sigma(t) = s'\}$; otherwise the equilibrium fails the refinement. But, if $\mu(s) = \sup\{t : \sigma(t) = s'\}$, then $\mathbb{E}_{t \sim \mu(s)}[t] > \mathbb{E}_{t \sim \mu(s')}[t]$; which implies that $r_\mu(s) > r_\mu(s')$ and hence, $q_{\sigma, \iota}(s) > q_{\sigma, \iota}(s')$. From increasing differences, then if t finds is at least as good in its equilibrium value, any $t' > t$ has a profitable deviation. Hence, (σ, ι, μ) cannot be an equilibrium. $\square$

D.2.1 Consequences of failure

The following lemmas identify structures that necessarily lead to failure of the refinement.

Non-monotone indices are incompatible with equilibrium-dominance. In equilibrium, attention must be weakly increasing.

Lemma 12. *Fix an equilibrium (σ, ι, μ) and let $r_\mu(s) = u - \frac{c-s}{\tau_\mu(s)}$. If there exist $s', s \in S$ with $s' > s$ but $r_\mu(s') < r_\mu(s)$, then the equilibrium fails equilibrium-dominance.*

Proof. It cannot be that both $s', s \in S_\sigma$, since [Theorem 1](#) already rules out such non-monotonicity on-path. If instead $s' \in S_\sigma$, then the type issuing s' has a profitable deviation, so again the equilibrium fails. The only relevant case is $s' \in S_\sigma^-$.

Under optimistic beliefs at s' , we have $r_\mu(s') > 0$.²⁴ In that case, a type $t \in \text{supp}(\mu(s'))$ deviating from s' to s strictly improves both rank (earlier inspection) and unit margin $t - ps$. Hence s' cannot be a weak best response for t , so equilibrium-dominance forbids $\mu(s')$ assigning positive mass to such types. Contradiction. $\square$

Equilibrium-dominance rules out types with negative margins at inspected subsidies.

Lemma 13. *Fix an equilibrium (σ, ι, μ) . If some s is inspected with positive probability, i.e. $q_{\sigma, \iota}(s) > 0$, then the equilibrium fails equilibrium-dominance if there exists $t < ps$ with $t \in \text{supp}(\mu(s))$. In particular, if $\tau_\mu(s) < ps$ whenever $q(s) > 0$, then the equilibrium fails equilibrium-dominance.*

Proof. If $t < ps$, then $t - ps < 0$. Such a type never weakly prefers any inspected s . Hence equilibrium-dominance rules out $\mu(s)$ placing mass on t . $\square$

Any interior pooling implies a failure of equilibrium-dominance.

²⁴We assume generically that $u > c$, which guarantees that even a zero subsidy is inspected under “type-1” beliefs. In the main text we already impose the slightly weaker condition $u > c - 1/p$.

Lemma 14. Fix an equilibrium (σ, ι, μ) . If there exists $s \in (0, c)$ such that $\Pr\{t \in T_\sigma(s)\} > 0$, then (σ, ι, μ) fails equilibrium-dominance.

Proof. Assume, for a contradiction, that (σ, ι, μ) does not fail equilibrium-dominance. By Lemma 12, the index $r_\mu(\cdot)$ and the induced attention $q_\sigma(\cdot)$ are weakly increasing on $[0, c]$.

Let $t^* = \sup T_\sigma(s)$ be the highest type pooled at s . Consider an increasing sequence $t^k \uparrow t^*$ with $t^k \in T_\sigma(s)$ for all k . By definition, for every $\epsilon > 0$, there exists $N(\epsilon)$ such that for all $k \geq N(\epsilon)$, $|t^k - t^*| < \epsilon$.

From Lemma 10 there exists some $\epsilon > 0$ such that $(s, s + \epsilon) \cap S_\sigma = \emptyset$. Then $s' \in (s, s + \epsilon)$ is off-path and there exists a fix $q' \in [0, 1]$ such that with optimistic beliefs at s' , $q' = q_{\sigma, \iota'}(s') > q_{\sigma, \iota'}(s)$.²⁵ Thus, for any t^k

$$(t^k - ps')q' - (t^k - ps)q(s) = (q' - q(s))(t^* - ps) - (q' - q(s))(t^* - t^k) - pq'(s' - s).$$

Since we are on an equilibrium, for $k > 1$, $t^k > ps$, and thus $\lim_k (t^k - ps) = \pi^* > 0$. Then, fix $\epsilon_1 < \pi^*/2$ and set $s' = s + \epsilon_2$ with $\epsilon_2 < (q' - q(s))\pi^*/(2(pq')) < 0$, if $k > N(\epsilon_1)$, then

$$(t^k - ps')q' - (t^k - ps)q(s) = (q' - q(s))\pi^* - (q' - q(s))\epsilon_1 - pq'\epsilon_2 > 0.$$

That is, type- t^k has a profitable deviation. Contradiction □

We next pin down the inspected portion of σ and its lower boundary. Let $\underline{t} = \frac{pc}{1+p}$ be the marginal type that is just worth inspecting and just willing to subsidize.

Lemma 15. Fix an equilibrium (σ, ι, μ) . If the lowest inspected type exceeds $\underline{t}$, ie $t_0 = \inf\{t \in T : q_\sigma(t) > 0\} > \underline{t}$, the equilibrium fails equilibrium-dominance.

Proof. Suppose inspection starts at some $t_0 > \underline{t}$. Let $s_0 = \sigma(t_0)$ be the lowest on-path subsidy with positive attention.

First, if $\pi_{\sigma, \iota}(s_0; t_0) = \pi_0 > 0$, then there exists some $t \in (\underline{t}, t_0)$ such that $(t - ps_0) > 0$ and its attention is $q(s_0) > 0$ by construction. Hence the deviation yields strictly positive profit for t , $\pi_{\sigma, \iota}(s_0, t) > 0 = \pi_{\sigma, \iota}(\sigma(t), t)$.

Second, if $\pi_0 = 0$, then type t_0 has a profitable deviation to $s \in (c/(1+p), s_0)$. This follows as $t_0 > pc/(1+p)$, thus for such set $(t_0 - ps) > 0$ and under optimistic beliefs at s , $q_{\sigma, \iota'}(s) > 0$. Thus, $\pi_{\sigma, \iota'}(s; t_0) > 0 = \pi_0$. Since $q_\sigma(s) = 0$, that means that there exists at least another type $t' \in \text{supp}(\mu(s))$ such that $t' < ps$. Otherwise, $r_\mu(s) > 0$ and hence inspected. But, t' has a negative margin at s and from Lemma 13, this equilibrium fails equilibrium-dominance. □

²⁵In particular, $q' = \lim_{t \downarrow t^*} q^{\text{sep}}(t)$ and $q' = 1$ if $\sigma(1) = s$. This follows as in equilibrium, a firm is inspected only when all the firms that rank better fail to match. That is what q^{sep} captures (see Appendix C.1).

D.3 Survival and uniqueness

Lemma 16. *The reasonable equilibrium survives the equilibrium–dominance refinement.*

Proof. Suppose, for a contradiction, that there exists an off–path message that violates equilibrium–dominance. Any such message must fall into one of two cases: (i) $s \in (0, \underline{t}/p)$, or (ii) $s \in (\lim_{\varepsilon \uparrow 0} \sigma^*(t_1 + \varepsilon), c)$ when there is a discontinuity at the top.

Consider case (i). All types $t \leq \underline{t}$ earn nonnegative profit in equilibrium. If any of them were to deviate to an off–path message in $(0, \underline{t}/p)$, their profit would be strictly negative, even under optimistic beliefs. Hence no profitable deviation exists.

Now consider case (ii). At the reasonable equilibrium, beliefs already assign probability one to the highest possible type at $s = \sigma^*(t_1)$. Thus optimistic beliefs at any $s \in (\sigma^*(t_1), c)$ cannot make the deviation more attractive. A type that raises its subsidy above $\sigma^*(t_1)$ obtains the same attention as t_1 but must pay strictly more, which strictly reduces profit. Lower–type firms do not gain either: although attention may rise, the higher subsidy cost dominates, again reducing profit.

In both cases, no type finds it profitable to deviate to an off–path subsidy. Therefore, the reasonable equilibrium survives the equilibrium–dominance refinement. $\square$

Therefore, we are ready to establish the proof for the main result of the section.

Proof of Theorem 2. By Lemma 14, the policy in any refined equilibrium is step–increasing–step: no inspection below some cutoff, then a strictly increasing revealing branch, then a single pooling block at c . By Lemma 15 the lower cutoff must be $\underline{t}$. By Lemma 11, the revealing branch is characterized by the same truth–telling ODE and the same boundary condition at $\underline{t}$, hence it coincides with σ^* on $[\underline{t}, \bar{t})$. By Lemma 13 pooling can occur only at c , and by Lemma 8 the pooling cutoff $\bar{t}$ is uniquely pinned by boundary indifference. Thus the outcome (cutoffs, revealing schedule, and top pooling) is unique and matches the one induced by the reasonable equilibrium (σ^*, q^*, μ^*) . $\square$

E Omitted proofs for Section 5

E.1 Some useful identities

First, we show how to derive the simplifications presented for the consumer surplus and the producer surplus.

Lemma 17. *The producer surplus can be simplified as follows:*

$$\frac{PS_n(c, p)}{n} = \int_{\underline{t}}^1 (1 - F(t)) q^*(t) dt = \mathbb{E}_t \left[\frac{1 - F(t)}{f(t)} q^*(t) \right], \quad (16)$$

while the expected firms transfers are given by

$$\begin{aligned}\phi_n^F(p, c) &= \int_{\underline{t}}^{\bar{t}} \left(t - \frac{1 - F(t)}{f(t)} \right) q^{\text{sep}}(t) f(t) dt + [1 - F(\bar{t})] (\bar{t} q^{\text{pool}}(\bar{t})) \\ &= \mathbb{E}_t \left[\psi(t) q^{\text{sep}}(t; \underline{t}) \mid t \leq \bar{t} \right] F(\bar{t}) + (\bar{t} q^{\text{pool}}(\bar{t})) [1 - F(\bar{t})]\end{aligned}\tag{17}$$

Proof. Recall that:

$$\begin{aligned}\phi_n^F(p, c) &= \mathbb{E}_t [p \sigma^*(t) q^*(t)] \\ &= \int_{\underline{t}(p, c)}^{\bar{t}(p, c)} \left(t q^{\text{sep}}(t) - \int_{\underline{t}(p, c)}^t q^{\text{sep}}(x) dx \right) f(t) dt + p c q^{\text{pool}}(\bar{t}(p, c)) [1 - F(\bar{t}(p, c))]\end{aligned}$$

This expression can be further simplified. First note that

$$\begin{aligned}\int_{\underline{t}(p)}^{\bar{t}(p)} \left(t q^{\text{sep}}(t) - \int_{\underline{t}(p)}^t q^{\text{sep}}(x) dx \right) f(t) dt \\ &= \int_{\underline{t}(p)}^{\bar{t}(p)} \left(t - \frac{F(\bar{t}(p)) - F(t)}{f(t)} \right) q^{\text{sep}}(t) f(t) dt \\ &= \int_{\underline{t}(p)}^{\bar{t}(p)} \left(t - \frac{1 - F(t)}{f(t)} \right) q^{\text{sep}}(t) f(t) dt + [1 - F(\bar{t}(p))] \int_{\underline{t}(p)}^{\bar{t}(p)} q^{\text{sep}}(t) dt \\ &= \int_{\underline{t}(p)}^{\bar{t}(p)} \left(t - \frac{1 - F(t)}{f(t)} \right) q^{\text{sep}}(t) f(t) dt + [1 - F(\bar{t}(p))] (\bar{t}(p) - p c) q^{\text{pool}}(\bar{t}(p)),\end{aligned}$$

where the first equality uses integration by parts, the second follows from algebraic manipulation and the third from the [Boundary indifference](#) condition. Therefore, by pasting this back to the original expression it yields

$$\begin{aligned}\phi_n^F(p, c) &= \int_{\underline{t}}^{\bar{t}} \left(t - \frac{1 - F(t)}{f(t)} \right) q^{\text{sep}}(t) f(t) dt + [1 - F(\bar{t})] (\bar{t} q^{\text{pool}}(\bar{t})) \\ &= \mathbb{E}_t \left[\psi(t) q^{\text{sep}}(t; \underline{t}) \mid t \leq \bar{t} \right] F(\bar{t}) + (\bar{t} q^{\text{pool}}(\bar{t})) [1 - F(\bar{t})]\end{aligned}$$

Similarly, we can simplify the expression for the producer surplus. Recall that

$$\begin{aligned}
PS(p, c) &= \mathbb{E}_t[(t - p\sigma^*(t))q^*(t)] \\
&= \int_{\underline{t}}^{\bar{t}} (t - p\sigma^{\text{sep}}(t))q^{\text{sep}}(t)f(t) dt + \int_{\bar{t}}^1 (t - pc)q^{\text{pool}}(\bar{t})f(t) dt \\
&= \int_{\underline{t}}^{\bar{t}} \left(\int_{\underline{t}}^t q^{\text{sep}}(x) dx \right) f(t) dt + q^{\text{pool}}(\bar{t}) \left(\int_{\bar{t}}^1 tf(t) dt - pc[1 - F(\bar{t})] \right) \\
&= \int_{\underline{t}}^{\bar{t}} [F(\bar{t}) - F(t)]q^{\text{sep}}(t) dt + q^{\text{pool}}(\bar{t}) \left(\int_{\bar{t}}^1 tf(t) dt - pc[1 - F(\bar{t})] \right)
\end{aligned}$$

where the first equality is definition, the second is achieved by substitution, the third uses the definition and algebraic manipulation and the fourth integration by parts. Re-arranging we obtain

$$PS_n(p, c) = \int_{\underline{t}}^{\bar{t}} [1 - F(t)]q^{\text{sep}}(t) dt + q^{\text{pool}}(\bar{t}) \int_{\bar{t}}^1 tf(t) dt - [1 - F(\bar{t})] \left(\int_{\underline{t}}^{\bar{t}} q^{\text{sep}}(t) dt + q^{\text{pool}}(\bar{t})pc \right)$$

To conclude, note that from the [Boundary indifference](#)

$$\int_{\underline{t}}^{\bar{t}} q^{\text{sep}}(t) dt + q^{\text{pool}}(\bar{t})pc = (\bar{t} - pc)q^{\text{pool}}(\bar{t}) + q^{\text{pool}}(\bar{t})pc = \bar{t}q^{\text{pool}}(\bar{t})$$

Thus,

$$\begin{aligned}
PS_n(p, c) &= \int_{\underline{t}}^{\bar{t}} [1 - F(t)]q^{\text{sep}}(t) dt + q^{\text{pool}}(\bar{t}) \left[\int_{\bar{t}}^1 tf(t) dt - [1 - F(\bar{t})]\bar{t} \right] \\
&= \int_{\underline{t}}^{\bar{t}} [1 - F(t)]q^{\text{sep}}(t) dt + q^{\text{pool}}(\bar{t}) \left[1 - \bar{t}F(\bar{t}) - \int_{\bar{t}}^1 F(t) dt - [1 - F(\bar{t})]\bar{t} \right] \\
&= \int_{\underline{t}}^{\bar{t}} [1 - F(t)]q^{\text{sep}}(t) dt + \int_{\bar{t}}^1 [1 - F(t)]q^{\text{pool}}(\bar{t}) dt
\end{aligned}$$

Therefore,

$$PS_n(c, p) = \int_{\underline{t}}^1 (1 - F(t))q^*(t) dt = \mathbb{E}_t \left[\frac{1 - F(t)}{f(t)} q^*(t) \right]$$

□

Lemma 18. Fix an equilibrium (σ, ι, μ) and let $t_0 = \inf\{t : q_{\sigma, \iota}(\sigma(t)) > 0\}$. The expected match probability is

$$m(\sigma, \iota) = 1 - \left(1 - \int_{t_0}^1 tf(t) dt \right)^n.$$

Proof. A match occurs if, and only if, one inspected firm succeeds. Hence, the probability of no match equals the probability that all eligible firms fail.

Let $K \sim \text{Bin}(n, 1 - F(t_0))$ be the number of firms that are positively inspected and $\tau := \mathbb{E}[t \mid t \geq t_0]$ the mean success probability conditional on inspection. Conditional on $K = k$, the k eligible firms are i.i.d., so the probability they all fail is $(1 - \tau)^k$. Averaging over K yields

$$\begin{aligned} \Pr\{\text{no match}\} &= \sum_{k=0}^n \binom{n}{k} (1 - F(t_0))^k (F(t_0))^{n-k} (1 - \tau)^k \\ &= \left(F(t_0) + (1 - F(t_0))(1 - \tau) \right)^n \\ &= \left(1 - \mathbb{E}[t \mid t \geq t_0] \Pr\{t \geq t_0\} \right)^n \end{aligned}$$

Therefore,

$$m(\sigma, \iota) := \Pr\{\text{match}\} = 1 - \left(1 - \int_{t_0}^1 t dF(t) \right)^n. \quad (18)$$

□

E.2 Omitted proofs for Section 5

Proof of Lemma 5. (i) *Lower cutoff.* The expression for the lower cutoff $\underline{t}(p, c) = pc/(1 + p)$ follows directly from the consumer's participation condition. It is strictly increasing in c and p , and by inspection does not depend on n .

(ii) *Effect of market size.* Recall that

$$q_n^{\text{sep}}(t; t_0) = \left(1 - \int_t^1 x f(x) dx \right)^{n-1} \mathbb{1}\{t \geq t_0\}.$$

Since the term in parentheses is less than one, $q_n^{\text{sep}}(t; t_0)$ is strictly decreasing in n for each $t \in [t_0, 1)$. Moreover, by construction $q_n^{\text{sep}}(t; \cdot)$ is decreasing in t . Moreover, note then that

$$\frac{\partial}{\partial c} q_n^{\text{sep}}(t; \underline{t}) = \frac{\partial}{\partial t_0} q_n^{\text{sep}}(t; \underline{t}) \frac{\partial}{\partial c} \underline{t} \leq 0$$

Similarly,

$$\frac{\partial}{\partial p} q_n^{\text{sep}}(t; \underline{t}) = \frac{\partial}{\partial t_0} q_n^{\text{sep}}(t; \underline{t}) \frac{\partial}{\partial p} \underline{t} \leq 0.$$

(iii) *Monotonicity of policies.* The separating subsidy policy is given by

$$\sigma_n^{\text{sep}}(t; t_0) = \frac{1}{p} \left(t - \frac{\int_{t_0}^t q_n^{\text{sep}}(x) dx}{q_n^{\text{sep}}(t)} \right).$$

It is immediate that if t_0 increases, then for any $t \geq t_0$ the numerator $\int_{t_0}^t q_n^{\text{sep}}(x) dx$ falls, and hence $\sigma_n^{\text{sep}}(t; t_0)$ rises. This proves monotonicity in the cutoff.

To compare across market sizes, note that

$$\sigma_{n+1}^{\text{sep}}(t; t_0) - \sigma_n^{\text{sep}}(t; t_0) = \frac{1}{p} \left(\frac{\int_{t_0}^t q_n^{\text{sep}}(x) dx}{q_n^{\text{sep}}(t)} - \frac{\int_{t_0}^t q_{n+1}^{\text{sep}}(x) dx}{q_{n+1}^{\text{sep}}(t)} \right).$$

Since

$$q_{n+1}^{\text{sep}}(t) = q_n^{\text{sep}}(t) \left(1 - \int_t^1 x f(x) dx \right),$$

and $q_n^{\text{sep}}(\cdot)$ is increasing in t , the difference above is strictly positive. Hence $\sigma_{n+1}^{\text{sep}}(t; t_0) > \sigma_n^{\text{sep}}(t; t_0)$ for all $t > t_0$.

To conclude, note that for $t > \underline{t}$,

$$\frac{\partial}{\partial c} \sigma_n^{\text{sep}}(t; \underline{t}) = \frac{\partial}{\partial t_0} \sigma_n^{\text{sep}}(t; \underline{t}) \frac{\partial}{\partial c} \underline{t} \geq 0$$

Moreover

$$\begin{aligned} \frac{\partial}{\partial p} \sigma^{\text{sep}}(t; \underline{t}) &= \frac{\partial}{\partial p} \left(\frac{t}{p} - \frac{\int_{\frac{cp}{1+p}}^t q^{\text{sep}}(x) dx}{p q^{\text{sep}}(t)} \right) \\ &= -\frac{t}{p^2} - \left(\frac{-q^{\text{sep}}(\underline{t})(c/(1+p)^2) - q^{\text{sep}}(t)^2(t - p\sigma^{\text{sep}}(t))}{p^2 q^{\text{sep}}(t)^2} \right) \\ &= -\frac{1}{p^2} \left(p\sigma^{\text{sep}}(t; \underline{t}) - \frac{q^{\text{sep}}(\underline{t})}{q^{\text{sep}}(t)^2} \frac{c}{(1+p)^2} \right) \end{aligned}$$

Moreover,

$$p\sigma^{\text{sep}}(t; \underline{t}) > \frac{q^{\text{sep}}(\underline{t})}{q^{\text{sep}}(t)^2} \frac{c}{(1+p)^2} \iff p\sigma^{\text{sep}}(t; \underline{t})q^{\text{sep}}(t) > \frac{c}{(1+p)^2} = \frac{\underline{t}}{p(1+p)}.$$

Therefore, if $p(1+p) \geq 1$, this last equality is true and thus $\frac{\partial}{\partial p} \sigma^{\text{sep}}(t; \underline{t}) < 0$. Otherwise, as $p(1+p) < 1$, then $-\underline{t} > -\frac{c}{(1+p)^2}$, and we have that

$$p\sigma^{\text{sep}}(t; \underline{t}) > \frac{q^{\text{sep}}(\underline{t})}{q^{\text{sep}}(t)^2} \underline{t} \implies p\sigma^{\text{sep}}(t; \underline{t}) > \frac{q^{\text{sep}}(\underline{t})}{q^{\text{sep}}(t)^2} \frac{c}{(1+p)^2}$$

We thus establish the result.

(iv) *Cutoff monotonicity.* As n increases, subsidies rise and the subsidy cap c is reached at a lower type. Similarly, an increase in p or c shifts the incentive constraints so that the upper

cutoff $\bar{t}$ is attained more quickly. Thus $\bar{t}$ is strictly decreasing in both c and p .

This completes the proof. $\square$

Proof of Theorem 3. Let $\bar{q}_n(p, c) := Q_n(p, c)/n$ be the per-firm inspection probability.

(i) *Search intensity.* Fix (p, c) and compare n and $n + 1$. Since $\bar{q}_n$ is the per-firm inspection probability, $\bar{q}_{n+1} \leq \bar{q}_n$ by Lemma 5(ii). However, total inspections

$$Q_n = n \bar{q}_n \quad \text{and} \quad Q_{n+1} = (n + 1) \bar{q}_{n+1}$$

satisfy $Q_{n+1} \geq Q_n$.

Therefore Q_n is (weakly) increasing in n . Next, holding n fixed and raising p or c increases $\bar{t}$ and weakly expands the top pool, so $q^*(t)$ weakly falls for every t ; hence Q_n is decreasing in p and c .

(ii) *Match probability.* As the price of subsidy or search cost increases, the number of participating types in equilibrium lowers (see Lemma 5). Thus, the match probability decreases. However, as the number of competing firms increases, the match probability increases as the failure event is lower. That means, match probability increases in n .

(iii) *Consumer surplus.* The consumer surplus is decreasing in p . This follows as the match probability decreases as p increases, and also firms transfers. A similar argument shows that as competition increases, consumer surplus also increases.

(iv) *Producer surplus.* From Lemma 17, it follows that

$$PS_n(c, p) = \mathbb{E}_t \left[\frac{1 - F(t)}{f(t)} q_n^*(t; p, c) \right]$$

where $q_n^*(t; p, c)$ is the attention function (equal to the separating one for $t \in [t_0, t_1)$ and equal to the pooling one in $t \in (t, 1]$).

Write $w(t) := 1 - F(t)$ and, for any attention schedule q , its cumulative $Q_n(x) := \int_0^x q_n^*(s; p, c) ds$. Integration by parts yields

$$\int_0^1 w(t) q(t) dt = w(1) Q(1) - \int_0^1 w'(t) Q(t) dt.$$

Since $w'(t) = -f(t) \leq 0$ and $w(1) = 1 - F(1) = 0$ (support $[0, 1]$), this reduces to

$$PS_n(c, p) = \int_0^1 w(t) q_n^*(t; p, c) dt = - \int_0^1 w'(t) Q_n(t; p, c) dt = \int_0^1 f(t) Q_n(t; p, c) dt.$$

(i) Fix n, c . If $p' > p$ then following Lemma 5, $Q_n(t; p', c) \leq Q_n(t; p, c)$ for all t . Because

$f(t) \geq 0$ for all t ,

$$PS_n(c, p') = \int_0^1 f(t) Q_n(t; p', c) dt \leq \int_0^1 f(t) Q_n(t; p, c) dt = PS_n(c, p).$$

(ii) The same argument applies to c : if $c' > c$ then $Q_n(t; p, c') \leq Q_n(t; p, c)$ for all t , hence $PS_n(c', p) \leq PS_n(c, p)$.

(iii) Fix p, c . Once again, $Q_{n+1}(t; p, c) \geq Q_n(t; p, c)$ for all t , then

$$PS_{n+1}(c, p) = \int_0^1 f(t) Q_{n+1}(t) dt \geq \int_0^1 f(t) Q_n(t) dt = PS_n(c, p),$$

so aggregate producer surplus is nondecreasing in n .

□

F Omitted proofs for [Section 6](#)

Proof of [Theorem 4](#). (i) If pooling were inactive at the maximizer (so $\bar{t} = 1$), a small decrease in p would create a positive-measure pool and increase the pool term in (13) at first order, while affecting the revelation term only at second order through the cutoff; revenue would rise, a contradiction.

(ii) Let t_ψ solve $\psi(t_\psi) = 0$ (exists and is unique under regularity). If $\bar{t}(p^*) \geq t_\psi$, then reducing p slightly lowers $\bar{t}$ and expands the pool into the region with $\psi < 0$. Using boundary indifference and the envelope logic behind (13), the marginal revenue gain from the pool dominates the marginal loss on the revelation term at that point, so revenue rises. Hence at an interior maximum we must have $\bar{t}(p^*) < t_\psi$, and some $\psi < 0$ types are inspected (in the pool or just below it).

□