

How to Segment a Search Market: Information Design and Directed Search

Teddy Mekonnen*

September 3, 2025

Abstract

This paper examines when the public provision of information in search markets improves welfare. I consider a two-sided frictional search market in which buyers match with vertically differentiated sellers. The market is segmented into submarkets based on seller types. Such segmentation serves as a public signal that buyers use to direct their search. Given a segmentation, I characterize both the socially efficient and the equilibrium allocation of buyers across submarkets, and identify a Hosios-type condition under which the equilibrium allocation is efficient. I then analyze the design of surplus-maximizing segmentations, showing that the nature of search externalities determines when the constrained-efficient segmentation fully separates seller types or pools them into at most a binary partition.

JEL Classifications: C78, D62, D83

Keywords: Directed search, market segmentation, information design

***Contact:** mekonnen@brown.edu. Department of Economics, Brown University.

I would like to express my deepest appreciation to Eddie Dekel, Federico Echenique, Bruno Strulovici, and Asher Wolinsky for their guidance and mentorship. This work has also greatly benefited from conversations with Alessandro Bonatti, Odilon Câmara, Hector Chade, Yeon-Koo Che, Pedro Dal Bó, Laura Doval, Jack Fanning, Kyungmin Kim, Stelios Michalopoulos, Bobak Pakzad-Hurson, João Ramos, Jesse Shapiro, Nikhil Vellodi, and Rajiv Vohra. Finally, I thank the editor and anonymous referees for comments and suggestions that greatly improved the paper. This paper supersedes a previous draft entitled “Random versus Directed Search for Scarce Resources.”

1. Introduction

Agents in search markets frequently rely on publicly available information to guide their search for desirable goods and services. For example, firms often recruit job candidates from reputable universities; shoppers on an online marketplace, such as *Amazon*, filter for highly rated products; and health insurance enrollees consult review aggregators like *Vitals* or *Healthgrades* to identify top-rated physicians. In each of these settings, agents use information about the quality of goods or services to segment the market into smaller submarkets before searching for a match in one of these submarkets. I refer to such information, which aids agents in directing their search prior to finding a match, as *ex-ante information*.

In single-agent search environments, the social value of information is unambiguous: more precise ex-ante information improves efficiency by facilitating matches between agents and higher-quality goods and services. In real-world search markets, however, many agents compete for limited resources. Universities graduate a finite number of students each year, online marketplaces have limited inventories of products, and insurance networks feature a restricted pool of physicians. Moreover, these resources are typically vertically differentiated—firms prefer more productive employees, shoppers value higher-quality products, and patients seek more competent physicians—with agents vying for higher-quality matches. In such environments, agents’ search decisions create externalities that influence the outcomes of others. Furthermore, these search externalities are shaped by the publicly available ex-ante information that guides the agents’ strategic decisions. Thus, a central question arises: does the public provision of more precise ex-ante information enhance welfare in search markets with externalities?

This paper addresses this question by considering a two-sided market populated by a finite measure of buyers and sellers, where each buyer (he) has a unit demand, and each seller (she) has a unit supply. The model has three defining features: First, sellers are vertically differentiated, with a seller’s private type determining the quality of the good she supplies. Second, buyers observe a public signal about seller types, which they use to guide their search. Third, search is frictional, with the likelihood of meeting a trading partner determined by the market tightness—the ratio of buyers to sellers. The meeting function, which governs this frictional search, encapsulates two key externalities: a thick market externality, where a tighter market facilitates more

meetings between buyers and sellers, and a congestion externality, where a tighter market intensifies competition among buyers, thereby reducing the probability that an individual buyer meets a seller. These opposing forces jointly determine equilibrium and welfare outcomes in search markets.

Public ex-ante information is modeled as a segmentation of the search market into a collection of submarkets. Sellers are allocated to these submarkets depending on their types, possibly with some randomization. I consider any flexible segmentation of the market à la [Bergemann et al. \(2015\)](#). Based on this segmentation, buyers form beliefs about the composition of seller types in each submarket. Each buyer then selects a submarket and engages in frictional search for a seller within it. If the buyer meets a seller, he observes her type, and the two bargain over the terms of trade.

The first set of results characterizes both the socially efficient and equilibrium allocations of buyers across submarkets for any fixed market segmentation. From a welfare perspective, tighter submarkets are advantageous because they facilitate more meetings, thereby increasing the overall volume of trade. However, with a finite measure of buyers, raising tightness in one submarket necessarily reduces it elsewhere. The unique socially efficient allocation therefore balances the thick market externality created by tighter submarkets against the expected surplus generated within each submarket.

In contrast, the unique equilibrium allocation arises from how buyers trade off the congestion in each submarket against the share of surplus they expect to capture in that submarket conditional on trade. In equilibrium, buyers follow a cutoff strategy based on a reservation value. They avoid submarkets where the expected payoff falls below this threshold. Among the remaining submarkets, those offering higher payoffs conditional on trade attract more buyers until their level of congestion precisely offsets their potential gains from trade.

Interestingly, a higher reservation value leads to lower buyer welfare. When buyers adopt a higher reservation value, they enter fewer submarkets, which increases the congestion within the submarkets they do join. This reduces the probability that any one buyer meets a seller, thereby lowering their average payoff.¹ In equilibrium, however, the reservation value itself must be consistent with buyers' average payoff, giving rise to a rational expectations condition.

¹This contrasts with single-agent search environments, where a higher reservation value does not generate congestion externalities and typically corresponds to a higher expected payoff.

This decentralized decision-making often leads to an inefficient allocation for two reasons: First, the seller types that generate the greatest total surplus may differ from those that provide buyers with the largest share of surplus, leading to a misalignment between the preferences of buyers and those of a social planner. Second, even when the preferences of the planner and the buyers coincide, buyers may over-concentrate in submarkets containing high-type sellers, creating excessive congestion. This misallocation highlights the dual role of information: while it helps buyers direct their search, it could also amplify the negative effects of congestion externalities.

Despite these two channels for inefficiencies, I identify a necessary and sufficient condition under which the socially efficient allocation of buyers is obtained as an equilibrium outcome. The condition, similar to [Hosios \(1990\)](#), relates the efficiency of decentralized search to the elasticity of the meeting function and the bargaining power of buyers. When satisfied, this condition ensures that buyers internalize the congestion externalities they impose, aligning their private incentives with the planner's objectives and yielding an efficient equilibrium outcome. However, this condition is satisfied only in knife-edge cases, implying that equilibrium outcomes are inefficient in most practical settings.

The second set of results examines the design of surplus-maximizing market segmentations under two settings: one where buyers' allocations across submarkets is socially efficient (the first-best segmentation) and another where buyers' allocations reflect equilibrium behavior (the constrained-efficient segmentation). The first-best segmentation reflects the planner's ability to directly control both the segmentation and the allocation of buyers, while the constrained-efficient segmentation accounts for buyers' strategic responses to the public provision of ex-ante information.

When buyers are efficiently allocated across submarkets, I show that a perfect market segmentation, where each seller type is assigned to a unique submarket, constitutes a first-best market segmentation. Intuitively, because more meetings lead to more trade, and because trade with a higher seller type generates a larger surplus, the socially efficient allocation of meetings must increase in seller types. Concentrating buyers into a single submarket for each seller type then allows the planner to achieve a desired volume of meetings while minimizing the number of buyers required. This segmentation enables the planner to tailor the allocation of meetings to seller types while fully leveraging the thick market externality.

Unlike the first-best setting, the constrained-efficient case requires the planner to

account for buyers’ strategic responses to the ex-ante information revealed by the market segmentation. This interdependence between segmentation design and equilibrium behavior complicates the characterization of the optimal segmentation. Nevertheless, under mild assumptions on the bargaining power of buyers, I show that the problem can be decomposed into two distinct steps: an information design problem, where the planner determines what information to reveal about seller types, and a fixed-point problem, which ensures consistency between the segmentation and the buyers’ equilibrium responses. This decomposition reveals that constrained-efficient market segmentations take on a monotone partition structure (Kleiner et al., 2021; Arieli et al., 2023).

Further refinement to the structure of constrained-efficient segmentations arises under additional assumptions regarding the nature of congestion externalities. If congestion externalities worsen gradually as market tightness increases, perfect segmentation constitutes a constrained-efficient segmentation. In this case, a fully directed search market obtains, where buyers have complete ex-ante information about seller types. In contrast, if congestion externalities worsen rapidly, the constrained-efficient segmentation takes the form of a binary partition. Sellers are grouped into two broad categories: seller types above a threshold are pooled into a “high” submarket, while types below the threshold are pooled into a “low” submarket. In this case, search is partially-directed with all buyers joining only the high submarket and engaging in random search within it.

These findings underscore the critical role of thick market and congestion externalities in determining the welfare consequences of more precise ex-ante information. When congestion externalities play a more limited role in search decisions, more informative public signals improve welfare, making perfect segmentation constrained-efficient. Conversely, when congestion externalities dominate, the constrained-efficient segmentation withholds substantial information from the buyers, demonstrating the potential negative welfare implications from the public provision of more precise ex-ante information in search markets with externalities.

1.1. Related Literature

The role of information in search markets can be broadly divided into two categories. The first is ex-ante information, where agents observe signals prior to meeting a potential match and use this information to guide their search strategies. The second

is interim information, where agents observe signals after meeting a potential match and use this information to determine the terms of trade.

This paper focuses on the former, and specifically, considers the role of ex-ante information in a market for search goods—goods whose quality is revealed upon inspection.² In this context, [Anderson and Renault \(2006\)](#), [Choi et al. \(2019\)](#), and [Lyu \(2023\)](#) analyze a monopolist’s optimal provision of information when buyers use it to decide between visiting the monopolist or taking their outside option. Similarly, [Choi et al. \(2018\)](#) explores the role of ex-ante information in oligopoly markets for horizontally differentiated search goods. However, these studies focus on settings with either a single buyer or sellers without capacity constraints, thereby abstracting away from search externalities. This paper addresses this gap by focusing on how ex-ante information interacts with thick market and congestion externalities in markets with multiple buyers and capacity-constrained sellers.³

This work also connects to the extensive literature on competitive search markets, where sellers post prices that buyers use to direct their search.⁴ Examples include search for exchange goods ([Butters, 1977](#); [Peters, 1991](#)), labor markets ([Montgomery, 1991](#); [Moen, 1997](#); [Mortensen and Wright, 2002](#)), and markets with two-sided heterogeneity ([Shi, 2002](#); [Shimer, 2005](#); [Eeckhout and Kircher, 2010a](#)). A central insight from this literature is that equilibrium outcomes are typically efficient because posted prices internalize search externalities.

In contrast, this paper considers directed search driven solely by information rather than posted prices. Terms of trade are instead determined via bargaining once a buyer and seller meet,⁵ as in [Diamond \(1982a\)](#), [Mortensen \(1982\)](#), and [Pissarides \(1985\)](#).

²In this paper, a buyer observes a seller’s type after the two meet, which reveals the quality of the seller’s product.

³Of course, other types of externalities could be considered. For example, [Vellodi \(2018\)](#) considers the role of ex-ante information in the entry and exit decisions of sellers in a search market, and shows that suppressing information can incentivize the entry of new sellers and delay the exit of incumbents.

⁴While the labor and matching literature often uses the terms *directed search* and *competitive search* interchangeably, I follow [Wright et al. \(2021\)](#) in distinguishing between the two. As per their definition, “Directed search means agents see some, although perhaps not all, characteristics of other agents, and based on that choose where to look for counterparties,” while “competitive search equilibrium means that agents on one side of the market post the terms of trade, while agents on the other side observe what is posted and direct their search accordingly.” Thus, competitive search is a special case of directed search.

⁵One such environment is the labor market for high-skilled workers: firms often direct their search by recruiting candidates from top universities, but salaries are negotiated after a successful completion of the interview process. Similarly, the posted-price framework is unsuitable in environments where

Consequently, prices cannot fully internalize the search externalities unless a Hosios condition (Hosios, 1990) is satisfied. Hence, this paper demonstrates that beyond price-based competitive search, directed search can give rise to inefficient equilibrium outcomes. Furthermore, it identifies conditions under which information design can mitigate these inefficiencies, offering insights into the optimal provision of ex-ante information in search markets.

Few papers explore directed search beyond price-based competitive search. The most relevant work in this regard is Menzio (2007), who examines a private-values labor market where informed firms send cheap-talk messages that uninformed workers use to direct their search. Once a trading partner is met, the terms of trade are determined via bargaining. Menzio (2007) shows that while informative cheap-talk equilibria exist, all equilibrium outcomes are inefficient.⁶ These inefficiencies stem from two interdependent externalities: search externalities on the workers' side of the market, and signaling externalities on the firms' side, which arise because the informational content of a firm's cheap-talk message depends on the messaging strategies of other firms.

In this paper, however, the market segmentation (and thus, its informational content) is taken as exogenously given by both buyers and sellers, reflecting the increasingly data-driven segmentation of markets based on certification or reviews rather than cheap talk. As a result, signaling externalities are absent from my model, leading to a clean analysis of how information interacts with search externalities. This, in turn, allows me to identify conditions under which efficient and equilibrium outcomes coincide.

Of course, directed search need not be based solely on posted prices or information; rather, it may be based on both. Eeckhout and Kircher (2010b) show that when buyers and sellers meet bilaterally, it is efficient to perfectly segment the market by type, and that this outcome can be decentralized by having buyers direct their search based on both the market segmentation and sellers' posted prices. Similarly, Cai et al. (2017) demonstrate that bilateral meetings are both sufficient and necessary for perfect market segmentations to emerge as socially efficient. In this respect, my

sellers lack commitment power, or where prices are fixed exogenously.

⁶When the trading mechanism is a first-price auction rather than bargaining, Kim and Kircher (2015) demonstrates the existence of a fully revealing and efficient cheap-talk equilibrium in a private-values setting, whereas Kim (2012) shows that no such equilibrium exists in a setting with vertical differentiation.

finding that perfect segmentation is efficient is not surprising. However, I show that this outcome can be implemented as an equilibrium if and only if the Hosios condition identified in the paper is satisfied.

Finally, this paper loosely relates to the body of work on the role of interim information, especially in markets for *experience goods*—goods whose quality is revealed only upon consumption. In the context of large search markets, [Lauermann \(2012\)](#) and [Lester et al. \(2019\)](#) demonstrate the negative welfare implications of interim information provision. A growing body of work also considers the design of interim information for experience goods in single-agent settings: [Romanyuk and Smolin \(2019\)](#), [Dogan and Hu \(2022\)](#), [Hu \(2022\)](#), [Mekonnen et al. \(2023\)](#), and [Sato and Shirakawa \(2023\)](#) all consider centralized provision of interim information, while [Board and Lu \(2018\)](#), [Au and Whitmeyer \(2023\)](#), [Mekonnen and Pakzad-Hurson \(2024\)](#), and [He and Li \(2023\)](#) consider competitive provision of information. Beyond search markets, the welfare implications and the design of interim information has also been analyzed in various other contexts, including price-discriminating monopoly ([Bergemann et al., 2015](#)), differentiated Bertrand competition ([Elliott et al., 2021](#)), bilateral matching markets [Condorelli and Szentes \(2023\)](#), and many-to-one matching markets [Li \(2023\)](#).

The remainder of the paper is structured as follows: I describe the model in Section 2, and derive the first-best outcome in Section 3. I then characterize the equilibrium outcome for a fixed market segmentation in Section 4. Finally, Section 5 discusses the efficiency properties of equilibrium outcomes and tackles the design of a constrained-efficient market segmentation. All proofs are in the [Appendix](#).

2. Model

A two-sided market is populated by a unit mass of homogeneous buyers on one side and a mass $k > 0$ of heterogeneous sellers on the other. Each buyer (he) has a unit demand for a good, and each seller (she) has a unit supply. A seller’s type, denoted by θ , captures the quality of her good. The type space Θ is a compact interval normalized to $[0, 1]$. Let $\mathcal{B}(\Theta)$ denote the Borel σ -algebra of Θ and $\Delta(\Theta)$ denote the space of probability measures over $(\Theta, \mathcal{B}(\Theta))$. Types are distributed according to an absolutely continuous probability measure $\mu \in \Delta(\Theta)$. Each seller knows her own type, but each buyer ex-ante knows only the type distribution; he observes a seller’s type only upon meeting her.

The market is segmented into submarkets, with the segmentation taken as exogenous and common knowledge among the buyers and sellers. Each seller is allocated to a submarket based on her type, possibly with some randomization. Each buyer, on the other hand, can freely choose which submarkets to join.

The timing is as follows: Given the market segmentation, each buyer picks a submarket to join. Within each submarket, buyers and sellers meet each other according to a *bilateral meeting function*, which will be formalized shortly. Upon meeting a seller, a buyer observes the seller's type, and the two bargain over the terms of trade. If they reach a mutual agreement to trade at some price $p \in \mathbb{R}$, the buyer and seller obtain payoffs $\theta - p$ and p , respectively. Buyer-seller pairs that fail to trade as well as buyers and sellers who do not find a trading partner earn a payoff of zero.

2.1. Market segmentation

A submarket is simply a market populated by a subset of seller types. Accordingly, each submarket can be indexed by the distribution of seller types it contains; that is, a submarket with seller-type distribution $\nu \in \Delta(\Theta)$ is referred to as submarket ν .

A market segmentation is an allocation of sellers to submarkets based on their types. Formally, let $\mathcal{B}(\Delta(\Theta))$ denote the Borel σ -algebra over $\Delta(\Theta)$, and let $\Delta(\Delta(\Theta))$ be the space of probability measures over $(\Delta(\Theta), \mathcal{B}(\Delta(\Theta)))$. A market segmentation is defined as a stochastic kernel $\sigma : \mathcal{B}(\Delta(\Theta)) \times \Theta \rightarrow [0, 1]$ such that,

- (a) For each $\theta \in \Theta$, the mapping $B \mapsto \sigma(B, \theta)$ is a probability measure in $\Delta(\Delta(\Theta))$;
- (b) For each $B \in \mathcal{B}(\Delta(\Theta))$, the mapping $\theta \mapsto \sigma(B, \theta)$ is $\mathcal{B}(\Theta)$ -measurable.

Under segmentation σ , a type- θ seller is allocated to submarket ν with probability $\sigma(d\nu, \theta)$. Consequently, out of the total mass k of sellers, the fraction allocated to submarket ν is given by

$$\sigma_\mu(d\nu) := \int_{\Theta} \sigma(d\nu, \theta) \mu(d\theta),$$

and the support $\text{supp}(\sigma_\mu)$ represents the submarkets present in the market.

Of course, the market segmentation must ensure that each submarket ν actually contains seller types distributed according to ν . In particular, σ must satisfy the following consistency condition: for all $A \in \mathcal{B}(\Theta)$ and $B \in \mathcal{B}(\Delta(\Theta))$,

$$\int_A \sigma(B, \theta) \mu(d\theta) = \int_B \nu(A) \sigma_\mu(d\nu). \quad (1)$$

This identity has a natural interpretation: consider sellers types in some measurable subset $A \subseteq \Theta$. The fraction of these types allocated to submarkets in B , given by the left-hand side, must equal the fraction of the same types contained in submarkets in B , given by the right-hand side.

The consistency condition in (1) guarantees that each submarket ν indeed has seller-type distribution ν , and that buyers can infer this using Bayes' rule. Thus, submarkets serve a dual role for buyers: each one is a platform for meeting a seller, and it also serves as a signal of the seller types therein. This is especially transparent when evaluating the consistency condition (1) at $B = \Delta(\Theta)$, which yields

$$\int_{\Delta(\Theta)} \nu \sigma_{\mu}(d\nu) = \mu.$$

This coincides with the definition of a market segmentation in [Bergemann et al. \(2015\)](#), and the Bayes-plausibility constraint in [Kamenica and Gentzkow \(2011\)](#).

2.2. Bilateral Meeting functions

Meetings are assumed to be bilateral, so a buyer meets at most one seller and vice versa. Following the literature on search and matching ([Petrongolo and Pissarides, 2001](#); [Rogerson et al., 2005](#)), the likelihood that a buyer and seller meet each other within a submarket depends only on the *submarket tightness*, which is the ratio of buyers to sellers in that submarket.

Formally, a submarket populated by a measure b of buyers and a measure s of sellers has a submarket tightness of $t := b/s \in \mathbb{R}_+ \cup \{\infty\}$. Each seller meets a buyer with probability $m(t)$ and meets no one with probability $1 - m(t)$, where the mapping $t \mapsto m(t)$ denotes the *meeting function*. As a measure $s \cdot m(t)$ of sellers meet buyers and meetings are bilateral, there must also be a measure $s \cdot m(t)$ of buyers that meet sellers. Thus, each buyer meets a seller with probability $m(t)/t$. The following assumptions on the meeting function are maintained for the remainder of the paper.

Assumption 1. *The meeting function satisfies the following:*

- (a) $m(t) \leq \min\{1, t\}$ for all $t \geq 0$, and
- (b) m is twice differentiable, strictly increasing, and strictly concave.

The first assumption implies that both $m(t)$ and $m(t)/t$ are well-defined probabilities. To motivate the second assumption, notice that as the submarket tightness increases, the market contains even more buyers relative to sellers, which has two effects: On the one hand, each seller in the submarket has more opportunities to meet a buyer, giving rise to a *thick market externality* on the sellers' side. This effect is captured by the monotonicity of the meeting function.⁷ On the other hand, each buyer in the submarket faces more competition to meet a seller, giving rise to a *congestion externality* on the buyers' side. This effect is captured by the strict concavity of the meeting function, which guarantees that the mapping $t \mapsto m(t)/t$ is strictly decreasing.

Let $\alpha := m(\infty)$ be a seller's highest probability of meeting a buyer, and let $\beta := m(0)/0$ be a buyer's highest probability of meeting a seller, with $\alpha, \beta \in (0, 1]$. In other words, an agent's probability of meeting a trading partner is highest in an extremely unbalanced submarket where the agent is on the scarce side. However, even in such extreme cases, the agent is not guaranteed a meeting if $\alpha < 1$ (for a seller) or $\beta < 1$ (for a buyer).

Examples of meeting functions that satisfy [Assumption 1](#) include the *CES meeting function*

$$m(t) = \frac{\alpha\beta t}{(\alpha^\rho + (\beta t)^\rho)^{1/\rho}}$$

with $\rho > 0$, and the *Urn-ball meeting function*

$$m(t) = \beta t \left(1 - \exp\left(\frac{-\alpha}{\beta t}\right) \right).$$

While [Assumption 1](#) is common in the search and matching literature, it is admittedly restrictive. However, by standard continuity arguments, the results of this paper extend to meeting functions that do not satisfy [Assumption 1](#), provided such functions can be approximated by meeting functions that do satisfy it. For example, consider the meeting function $m(t) = \min\{\alpha, \beta t\}$, which generalizes the *frictionless meeting function* given by $m(t) = \min\{1, t\}$. This meeting function does not satisfy Point (b) of

⁷Some earlier work in the matching and search literature uses the term “thick market externality” to refer to increasing returns to scale in the meeting function, which can give rise to multiple equilibrium outcomes in search markets ([Diamond, 1982b](#)). In contrast, and consistent with much of the recent literature, I assume constant returns to scale. See [Petrongolo and Pissarides \(2001\)](#) for a more in-depth survey.

[Assumption 1](#), as it is only weakly increasing, weakly concave, and non-differentiable at $t = \alpha/\beta$. Nevertheless, the CES meeting function approximates the frictionless one as $\rho \rightarrow \infty$, and thus, the insights of this paper extend to this class of generalized frictionless meeting functions.

2.3. Trade decisions

I assume that a buyer and seller negotiate the terms of trade only after they meet, and that their negotiations result in all ex-post efficient trades being realized. Specifically, suppose a buyer has met a type- θ seller. If they trade, they generate a total surplus of $\theta \in [0, 1]$; If they do not, the surplus is zero. Thus, they mutually agree to trade at some price $p(\theta)$, which requires $\theta \geq p(\theta)$ for the buyer to agree and $p(\theta) \geq 0$ for the seller to agree. In other words, the price $p(\theta)$ is some convex combination of the surplus θ and the zero outside option.

Furthermore, I assume that buyers capture a non-negligible share of the surplus from trading with some seller types; otherwise, the buyers would have no incentives to direct their search to any specific submarket, making the problem at hand trivial. Formally, I assume:

Assumption 2. *There exists a $\mathcal{B}(\Theta)$ -measurable function $\lambda : \Theta \rightarrow [0, 1]$ such that*

(a) **Ex-post efficiency:** $p(\theta) = (1 - \lambda(\theta))\theta$ for all $\theta \in \Theta$, and

(b) **Non-triviality:** $\mu(\{\theta \in \Theta : \lambda(\theta) > 0\}) > 0$.

I refer to λ as the *surplus-splitting function*, with $\theta - p(\theta) = \lambda(\theta)\theta$ representing the share of surplus captured by a buyer when he trades with a type- θ seller.

One foundation for $p(\cdot)$ is as the outcome of a Rubinstein bargaining problem: upon meeting, a buyer and a seller enter a “bargaining phase” in which the pair make alternating offers on how to divide the surplus. When buyers have a discount factor of $\varrho_b \in (0, 1)$ and type- θ sellers have a discount factor of $\varrho_\theta \in (0, 1)$, the unique solution to the bargaining phase yields a surplus-splitting function of $\lambda(\theta) = (1 - \varrho_\theta)/(1 - \varrho_b\varrho_\theta)$ when buyers make the first proposal, or $\lambda(\theta) = (1 - \varrho_b)/(1 - \varrho_b\varrho_\theta)$ when sellers make the first proposal.

The price $p(\theta)$ can also be rationalized as the solution to a generalized Nash-bargaining problem

$$\max_{p \in \mathbb{R}} (\theta - p)^{\lambda(\theta)} p^{1-\lambda(\theta)},$$

with $\lambda(\theta)$ capturing the buyer's bargaining power when he meets a type- θ seller.

Finally, $p(\theta)$ can also reflect the expected price in a setting where a buyer who, upon meeting a type- θ seller, makes a take-it-or-leave-it offer with probability $\lambda(\theta)$, while the seller does the same with the complementary probability. If $\lambda(\theta) = 1$ for all $\theta \in \Theta$, the model becomes analogous to one with non-transferable utility because a buyer can trade with any seller he meets at a price of zero.

3. First-best market segmentation

As a benchmark, consider a planner that seeks to maximize total surplus by choosing the market segmentation as well as the allocation of buyers across submarkets. A market segmentation σ determines the allocation of sellers, with $k\sigma_\mu(B)$ representing the mass of sellers allocated to submarkets in $B \in \mathcal{B}(\Delta(\Theta))$. Let $Q \in \Delta(\Delta(\Theta))$ be the buyer-allocation policy, with $Q(B)$ representing the mass of buyers allocated to submarkets in $B \in \mathcal{B}(\Delta(\Theta))$.

Buyers are allocated only to submarkets that contain sellers, so Q must be absolutely continuous with respect to σ_μ . In particular, there exists a measurable function $\tau : \Delta(\Theta) \rightarrow \mathbb{R}_+$ such that $Q(d\nu) = k\tau(\nu)\sigma_\mu(d\nu)$ for σ_μ -almost all $\nu \in \Delta(\Theta)$. In other words, the mapping $\nu \mapsto \tau(\nu)$ represents the ratio of buyers to sellers within each submarket and, thus, I refer to it as the *submarket tightness function*.

The market segmentation and the submarket tightness function together fully determine the buyer-allocation policy. Hence, given a market segmentation σ , choosing a buyer-allocation policy Q is equivalent to choosing a submarket tightness function τ that satisfies the following feasibility constraint:

$$k \int_{\Delta(\Theta)} \tau(\nu)\sigma_\mu(d\nu) = 1. \quad (2)$$

Given market segmentation σ and submarket tightness function τ , there is a measure $k\sigma_\mu(d\nu)m(\tau(\nu))$ of meetings in submarket $\nu \in \text{supp}(\sigma_\mu)$, with each meeting generating an expected surplus of $\mathbb{E}_\nu[\theta] := \int_\Theta \theta\nu(d\theta)$. Therefore, the first-best outcome is the solution to the following surplus-maximization problem:

$$\max_{\sigma, \tau} k \int_{\Delta(\Theta)} m(\tau(\nu))\mathbb{E}_\nu[\theta]\sigma_\mu(d\nu) \quad \text{s.t.} \quad (2). \quad (\text{FB})$$

To analyze this surplus-maximization problem, I decompose it into two distinct optimization steps: First, for each market segmentation, I characterize the surplus-maximizing submarket tightness function. Second, I determine the optimal market segmentation. To that end, given market segmentation σ , the surplus-maximizing submarket tightness function is a solution to

$$\max_{\tau} k \int_{\Delta(\Theta)} m(\tau(\nu)) \mathbb{E}_{\nu}[\theta] \sigma_{\mu}(d\nu) \quad \text{s.t.} \quad (2). \quad (\text{FB-}\sigma)$$

To characterize the solution to (FB- σ), let the inverse of the mapping $t \mapsto 1/m'(t)$ be denoted by f so that $f(y) = t$ if $1/m'(t) = y$. Since $m'(t)$ is continuous and strictly decreasing in t , $f(y)$ is continuous and strictly increasing in y . Furthermore, under [Assumption 1](#), $\lim_{t \rightarrow 0} m'(t) = \beta$ and $\lim_{t \rightarrow \infty} m'(t) = 0$,⁸ which implies that $f : [1/\beta, \infty) \rightarrow \mathbb{R}_+$ with $f(1/\beta) = 0$ and $\lim_{y \rightarrow \infty} f(y) = \infty$.

Proposition 1. *Fix a market segmentation σ . Let η_{σ} be the unique value of $\eta \in (0, \beta)$ that solves*

$$k \int_{\Delta(\Theta)} f\left(\frac{\mathbb{E}_{\nu}[\theta]}{\eta}\right) \cdot \mathbb{1}_{[\beta \mathbb{E}_{\nu}[\theta] > \eta]} \sigma_{\mu}(d\nu) = 1.$$

Then the submarket tightness function τ_{σ}^{FB} given by

$$\tau_{\sigma}^{FB}(\nu) = \begin{cases} 0 & \text{if } \beta \mathbb{E}_{\nu}[\theta] \leq \eta_{\sigma} \\ f\left(\frac{\mathbb{E}_{\nu}[\theta]}{\eta_{\sigma}}\right) & \text{if } \beta \mathbb{E}_{\nu}[\theta] > \eta_{\sigma} \end{cases}$$

solves (FB- σ). Furthermore, any other submarket tightness function τ that solves (FB- σ) satisfies $\tau = \tau^{FB}$ for σ_{μ} -almost everywhere.

From a social welfare perspective, thicker submarkets are desirable because they facilitate more meetings, and ultimately, more trade. However, with a finite measure of buyers, increasing the thickness of one submarket necessarily reduces the thickness of others. This trade-off produces two key implications for the first-best allocation of buyers across submarkets.

⁸Since m is strictly concave and $m(0) = 0$ ([Assumption 1](#)), $m'(t) < m(t)/t$ for all $t > 0$. Taking limits on both sides as $t \rightarrow 0$, we have $m'(0) \leq \beta$. Similarly, given $t > 0$, we have $m'(t) > (m(t') - m(t))/(t' - t)$ for all $t' > t$. Taking double limits on both side, first as $t \rightarrow 0$ and then as $t' \rightarrow 0$, yields that $m'(0) \geq \beta$.

First, some submarkets may be *inactive*—no buyers are allocated to them. Specifically, given a market segmentation σ , submarkets generating an expected surplus lower than η_σ/β are inactive, while those generating a strictly larger expected surplus are *active*. In other words, the planner focuses the limited number of buyers on submarkets with a “high-enough” expected surplus.

Second, the marginal value of an additional buyer in any one of the active submarkets must be equal. As the meeting function is concave, this condition requires that an active submarket’s tightness to be strictly increasing in its expected surplus. Furthermore, the marginal value of an additional buyer in any active submarket must exceed that in any inactive submarket. Otherwise, reallocating some buyers from an active to an inactive submarket would increase total surplus. Under this interpretation, any submarket with an expected surplus equal to η_σ/β is a marginal submarket that leaves the planner indifferent between keeping it inactive or allocating buyers to it.

Let us now address the task of characterizing the first-best market segmentation: what market segmentation, along with the corresponding surplus-maximizing submarket tightness function, should the planner choose? In other words, the first-best market segmentation is the solution to the following:

$$\max_{\sigma} k \int_{\Delta(\Theta)} m(\tau_{\sigma}^{FB}(\nu)) \mathbb{E}_{\nu}[\theta] \sigma_{\mu}(d\nu).$$

To characterize the first-best market segmentation, I begin by defining the notion of a *perfect market segmentation*, which allocates each seller type to a distinct submarket. Formally, for each $\theta \in \Theta$, let $\delta_{\theta} \in \Delta(\Theta)$ denote the Dirac measure at θ , so that for any $A \in \mathcal{B}(\Theta)$, $\delta_{\theta}(A) = \mathbb{1}_{[\theta \in A]}$. The perfect market segmentation is then defined by the stochastic kernel σ^{PS} such that, for all $B \in \mathcal{B}(\Delta(\Theta))$ and $\theta \in \Theta$, $\sigma^{PS}(B, \theta) = \mathbb{1}_{[\delta_{\theta} \in B]}$.

Proposition 2. *The perfect segmentation σ^{PS} is a first-best market segmentation.*

Intuitively, by choosing a market segmentation σ and a corresponding submarket tightness function τ_{σ}^{FB} , the planner is effectively allocating

$$k\mu(d\theta) \int_{\Delta(\Theta)} m(\tau_{\sigma}^{FB}(\nu)) \sigma(d\nu, \theta)$$

meetings to type- θ sellers. The concavity in the meeting function then implies that, for each seller type, concentrating buyers into a single thick submarket minimizes the mass of buyers needed to attain a desired volume of meetings. Consequently, perfect segmentation emerges as a first-best market segmentation.

Nevertheless, perfect segmentation is not the only first-best market segmentation. For instance, under perfect segmentation, we know from [Proposition 1](#) that the submarket for any seller type below $\eta_{\sigma PS}/\beta$ is inactive, while the submarket for any seller type strictly above $\eta_{\sigma PS}/\beta$ is active. Thus, a lower censorship market segmentation, which pools all seller types $\theta \leq \eta_{\sigma PS}/\beta$ into a single submarket and maintains a perfect segmentation for all types $\theta > \eta_{\sigma PS}/\beta$, would also constitute a first-best market segmentation.

4. Search Equilibrium

Let us now consider a search market in which each buyer strategically chooses which submarkets to enter. What is the equilibrium outcome of such a decentralized search environment?

To address this, consider an arbitrary market segmentation σ , and suppose buyers anticipate a submarket tightness function $\tau : \Delta(\Theta) \rightarrow \mathbb{R}_+$. Given (σ, τ) , a buyer who joins submarket $\nu \in \text{supp}(\sigma_\mu)$ expects to meet a seller with probability $m(\tau(\nu))/\tau(\nu)$. If the buyer meets a type- θ seller within the submarket (with $\theta \in \text{supp}(\nu)$), he further expects to earn a payoff $\lambda(\theta)\theta$ from trade. Therefore, the buyer's expected payoff from joining submarket ν is given by

$$U(\nu; \tau) := \frac{m(\tau(\nu))}{\tau(\nu)} \underbrace{\int_{\Theta} \lambda(\theta)\theta \nu(d\theta)}_{:=\mathbb{E}_\nu[\lambda(\theta)\theta]}.$$

Definition 1. *A search equilibrium of a market segmentation σ is given by a submarket tightness function $\tau : \Delta(\Theta) \rightarrow \mathbb{R}_+$ and a reservation value $u \in \mathbb{R}$ such that*

- (a) τ satisfies the feasibility constraint (2), and
- (b) (τ, u) jointly satisfy

$$U(\nu; \tau) \leq u \tag{3}$$

for σ_μ -almost all ν with equality if $\tau(\nu) > 0$.

In essence, buyers form expectations about submarket tightness τ and their reservation value u , which represents the average payoff they expect to earn from participating in the search market. In equilibrium, these expectations must be consistent with a feasible allocation of buyers across submarkets, as captured by (2), as well as with an optimal buyers' search strategy, as captured by (3). Intuitively, the latter condition states that submarkets that yield a payoff lower than u remain inactive, while those offering payoffs at least as large as u are active. Furthermore, among active submarkets, buyers must be indifferent between any two; otherwise, they would all gravitate toward the submarket offering the higher expected payoff.

Finally, a search equilibrium (τ, u) must satisfy a rational expectations condition: the anticipated reservation value u should be consistent with the buyers' actual ex-ante payoff. To see this, observe that a pair (τ, u) that satisfies Definition 1 implies that $m(\tau(\nu))\mathbb{E}_\nu[\lambda(\theta)\theta] = u\tau(\nu)$ for σ_μ -almost all ν from (3). Integrating over all submarkets then yields

$$\begin{aligned} \int_{\Delta(\Theta)} m(\tau(\nu))\mathbb{E}_\nu[\lambda(\theta)\theta]\sigma_\mu(d\nu) &= u \int_{\Delta(\Theta)} \tau(\nu)\sigma_\mu(d\nu) \\ &= \frac{u}{k}, \end{aligned}$$

where the last equality follows from (2). On the other hand, buyers' ex-ante payoff from an allocation policy $Q \in \Delta(\Delta(\Theta))$ with $dQ/d\sigma_\mu = k \cdot \tau$ is given by⁹

$$\int_{\Delta(\Theta)} U(\nu; \tau)Q(d\nu) = k \int_{\Delta(\Theta)} m(\tau(\nu))\mathbb{E}_\nu[\lambda(\theta)\theta]\sigma_\mu(d\nu).$$

Putting these two expressions together then yields the desired rational expectations condition:

$$u = k \int_{\Delta(\Theta)} m(\tau(\nu))\mathbb{E}_\nu[\lambda(\theta)\theta]\sigma_\mu(d\nu), \quad (4)$$

that is, the buyers' reservation value is equal to the share of the ex-ante surplus they capture in equilibrium.

To characterize the equilibrium, let the inverse of the mapping $t \mapsto t/m(t)$ be

⁹Recall that given segmentation σ , submarket $\nu \in \text{supp}(\sigma_\mu)$ has a measure $k \cdot \sigma_\mu(d\nu)$ sellers. Hence, if the submarket has tightness $\tau(d\nu)$, then it must have a measure $Q(d\nu)$ of buyers, where $Q(d\nu) = k \cdot \tau(d\nu) \cdot \sigma_\mu(d\nu)$.

denoted by g so that $g(y) = t$ if $t/m(t) = y$. Since $m(t)/t$ is continuous and strictly decreasing in t , $g(y)$ is continuous and strictly increasing in y . Furthermore, $m(0)/0 = \beta$ and $\lim_{t \rightarrow \infty} m(t)/t = 0$, which implies that $g : [1/\beta, \infty) \rightarrow \mathbb{R}_+$ with $g(1/\beta) = 0$ and $\lim_{y \rightarrow \infty} g(y) = \infty$.

Proposition 3. *Fix a market segmentation σ . There exists a search equilibrium $(\tau_\sigma^*, u_\sigma^*)$ such that $0 < u_\sigma^* < \beta \sup_{\theta \in \Theta} \lambda(\theta)\theta$, and*

$$\tau_\sigma^*(\nu) = \begin{cases} 0 & \text{if } \beta \mathbb{E}_\nu[\lambda(\theta)\theta] \leq u_\sigma^* \\ g\left(\frac{\mathbb{E}_\nu[\lambda(\theta)\theta]}{u_\sigma^*}\right) & \text{if } \beta \mathbb{E}_\nu[\lambda(\theta)\theta] > u_\sigma^* \end{cases}.$$

Furthermore, the equilibrium is essentially unique: any other search equilibrium (τ, u) satisfies $\tau = \tau_\sigma^*$ for σ_μ -almost everywhere, and $u = u_\sigma^*$.

To build intuition, consider a fixed reservation value u . When buyers anticipate a payoff of u , they optimally avoid entering any submarket ν satisfying $\beta \cdot \mathbb{E}_\nu[\lambda(\theta)\theta] \leq u$, even if these submarkets offer the highest probability of meeting a seller. This is because the expected payoff from trade in such submarkets is too low to justify participation. Among the remaining submarkets, those offering a higher payoff conditional on trade attract more buyers, leading to greater congestion. As a result, congestion in each active submarket precisely offsets its potential gains from trade. Consequently, the optimal submarket tightness for a given reservation value u is determined by the mapping

$$\nu \mapsto g\left(\frac{\mathbb{E}_\nu[\lambda(\theta)\theta]}{u}\right) \mathbb{1}_{[\beta \mathbb{E}_\nu[\lambda(\theta)\theta] > u]}.$$

Moreover, a buyer's average payoff across all active submarkets is given by

$$\int_{\Delta(\Theta)} m\left(g\left(\frac{\mathbb{E}_\nu[\lambda(\theta)\theta]}{u}\right)\right) \mathbb{1}_{[\beta \mathbb{E}_\nu[\lambda(\theta)\theta] > u]} \mathbb{E}_\nu[\lambda(\theta)\theta] \sigma_\mu(d\nu),$$

which is strictly decreasing in u . In other words, when all buyers anticipate a higher reservation value, fewer submarkets remain active, increasing congestion within those submarkets. This, in turn, lowers the probability that any given buyer meets a seller, thereby reducing their average payoff. Of course, in equilibrium, the reservation value u_σ^* is not arbitrary; rather, it is determined by the rational expectations condition,

captured in (4), which requires it to equal the buyer’s average payoff across all active submarkets.

A comparison of the socially efficient allocation of buyers across submarkets ([Proposition 1](#)) to the equilibrium allocation ([Proposition 3](#)) highlights two main sources of inefficiency. First, a social planner’s preference over seller types is strictly monotonic, as a higher type generates a larger surplus. In contrast, a buyer’s valuation of seller types depends on the share of surplus they capture, which may be non-monotonic in seller types depending on the surplus-splitting function. For example, if $\lambda(\theta) = 1 - \theta$, buyers prefer to meet intermediate-type sellers over low- and high-type sellers. This divergence in preferences over seller types can lead to a misalignment in how the planner and buyers evaluate submarkets. Consequently, the set of active submarkets in a search equilibrium may differ from that of the socially efficient allocation.

Second, even when the planner’s and buyers’ preferences over seller types align, the allocation of buyers across these submarkets may differ. In equilibrium, buyers allocate themselves to equalize their expected payoffs across active submarkets. In contrast, the planner allocates buyers to equalize the marginal value of an additional meeting across active submarkets. The equilibrium allocation is driven by the level of congestion in a submarket, $m(t)/t$, while the planner’s allocation is driven by the thick market externality, $m'(t)$. From [Assumption 1](#), $m(t)/t \geq m'(t)$ for all $t \geq 0$, implying that, even when a submarket is active under both the socially efficient and equilibrium allocations, too many buyers join the submarket in equilibrium, resulting in inefficiently high levels of congestion.

5. Efficiency of Search Equilibrium

5.1. Hosios condition

Despite the potential inefficiencies in the equilibrium allocation of buyers, this section identifies necessary and sufficient conditions under which the first-best outcome, as established in [Proposition 1](#) and [Proposition 2](#), can be decentralized. To state the result, let

$$\varepsilon(t) := \frac{m'(t)}{m(t)} t$$

denote the elasticity of the meeting function for $t \geq 0$.

Proposition 4. *Suppose the surplus-splitting function $\lambda : \Theta \rightarrow [0, 1]$ is continuous. The pair $(\sigma^{PS}, \tau_{\sigma^{PS}}^*)$ solves (FB) if and only if for all $\theta \leq \eta_{\sigma^{PS}}/\beta$,*

$$\lambda(\theta)\theta \leq \lambda\left(\frac{\eta_{\sigma^{PS}}}{\beta}\right)\frac{\eta_{\sigma^{PS}}}{\beta},$$

and for all $\theta > \eta_{\sigma^{PS}}/\beta$

$$\lambda(\theta) = \lambda\left(\frac{\eta_{\sigma^{PS}}}{\beta}\right)\varepsilon\left(f\left(\frac{\theta}{\eta_{\sigma^{PS}}}\right)\right).$$

The proposition relates the efficiency of decentralized search to the elasticity of the meeting function and the buyers' bargaining power, giving rise to a *Hosios condition* (Hosios, 1990). Intuitively, the condition ensures that the surplus captured by a buyer within any inactive submarket is no greater than what he could capture from joining the submarket of a type- $\eta_{\sigma^{PS}}/\beta$ seller (the marginal submarket in the first-best outcome). Moreover, a buyer's expected payoff from joining any active submarket is exactly balanced against the marginal value of an additional buyer in that submarket. This balance implies that buyers internalize the congestion and thick market externalities they create when choosing a submarket, aligning their private incentives with the social planner's objectives and leading to an efficient outcome.

An instructive special case that satisfies both conditions of Proposition 4 is when $\lambda(\theta) = 0$ for all $\theta \in \Theta$. In this case, buyers would be completely indifferent across all submarkets since they capture no surplus, eliminating any tension between private and social incentives. Clearly, efficiency is straightforward to attain in equilibrium under these conditions. However, this trivial case is ruled out by Assumption 2, which ensures that buyers capture a strictly positive share of surplus.

5.2. Constrained-efficient market segmentation

A key takeaway from Proposition 4 is the restrictive nature of the Hosios condition. Achieving efficiency in a decentralized search equilibrium often requires a knife-edge condition that rarely holds in many familiar settings. For instance, inefficiency persists when the surplus-splitting function is constant and positive, which is a natural assumption if seller types determine surplus without directly affecting bargaining power. This rarity of efficiency in equilibrium raises an important question: what is

the constrained-efficient segmentation of a search market? That is, how should the social planner optimally segment the market when it can no longer directly control the allocation of buyers across submarkets but can still influence their choices indirectly through the information they observe prior to selecting a submarket?

Formally, each market segmentation σ induces an essentially-unique search equilibrium $(\tau_\sigma^*, u_\sigma^*)$, as characterized by [Proposition 3](#). A planner who implements a segmentation σ expects to create a measure $k\sigma_\mu(d\nu)m(\tau_\sigma^*(\nu))$ of meetings in submarket $\nu \in \text{supp}(\sigma_\mu)$, with each meeting generating an expected surplus of $\mathbb{E}_\nu[\theta]$. Thus, the constrained-efficient market segmentation is the solution to the planner's second-best problem:

$$\max_{\sigma} k \int_{\Delta(\Theta)} m(\tau_\sigma^*(\nu)) \mathbb{E}_\nu[\theta] \sigma_\mu(d\nu). \quad (\text{SB})$$

This optimization problem is inherently complex: it involves optimizing over an infinite-dimensional choice variable that enters the objective function in a highly non-linear way. Specifically, the choice of segmentation influences buyers' beliefs about seller types within each submarket, which in turn affects how buyers allocate themselves across submarkets.¹⁰ In [Appendix A.2](#), I derive a sufficient and necessary condition that decouples these two interdependent effects. However, even this requires solving a constrained concavification problem ([Kamenica and Gentzkow, 2011](#)) with a continuum of states. As such, a general characterization of the constrained-efficient market segmentation appears intractable in many cases.

To make further progress, I introduce the following additional assumption:

Assumption 3. *For all $\theta \in \Theta$, $\lambda(\theta) = \ell$ for some $\ell \in (0, 1]$.*

Under this assumption, buyers capture a fixed share ℓ of surplus generated in any submarket. Hence, maximizing total surplus is now equivalent to maximizing buyers' welfare. Moreover, [Assumption 3](#) implies that, in equilibrium, any two submarkets with identical posterior means have the same submarket tightness and generate the same surplus, i.e., for $\nu, \nu' \in \text{supp}(\sigma_\mu)$ with $\mathbb{E}_\nu[\theta] = \mathbb{E}_{\nu'}[\theta]$, we have $\tau_\sigma^*(\nu) = \tau_\sigma^*(\nu')$, and $m(\tau_\sigma^*(\nu))\mathbb{E}_\nu[\theta] = m(\tau_\sigma^*(\nu'))\mathbb{E}_{\nu'}[\theta]$. This property allows us to reformulate [\(SB\)](#) as an optimization problem over posterior-mean distributions with a majorization constraint.

¹⁰Although [\(SB\)](#) resembles an information design problem, it differs in a crucial way: whereas a receiver's optimal action in standard information design depends solely on the realized posterior beliefs, buyers' search strategies here depend both on the entire distribution of posterior beliefs.

To that end, let F denote the cumulative distribution function (CDF) associated with the prior $\mu \in \Delta(\Theta)$. Importantly, F is also the posterior-mean distribution induced by the perfect market segmentation.¹¹ More generally, a CDF $H : \Theta \rightarrow [0, 1]$ represents a posterior-mean distribution induced by some market segmentation if and only if H is a mean-preserving contraction of F (Blackwell, 1953; Blackwell and Girshick, 1954). Let $\text{MPC}(F)$ denote the set of all such mean-preserving contractions.

For each $H \in \text{MPC}(F)$, there exists some market segmentation σ that induces H as its posterior-mean distribution, and yields an essentially-unique search equilibrium $(\tau_\sigma^*, u_\sigma^*)$. With some abuse of language, I refer to H as the market segmentation, and denote the search equilibrium as (τ_H^*, u_H^*) where $\tau_H^*(\mathbb{E}_\nu[\theta]) := \tau_\sigma^*(\nu)$ and $u_H^* := u_\sigma^*$.¹²

The planner's second-best problem (SB) can now be reformulated as follows:

$$\max_{H \in \text{MPC}(F)} k \int_{\Theta} m(\tau_H^*(x)) x dH(x). \quad (\text{SB}')$$

While this optimization problem remains non-linear, the next result provides a tractable characterization of a constrained-efficient market segmentation.

Proposition 5. *A segmentation $H \in \text{MPC}(F)$ is constrained-efficient if and only if*

$$H \in \arg \max_{\hat{H} \in \text{MPC}(F)} k \int_{\Theta} \tau_H^*(x) d\hat{H}(x). \quad (5)$$

Furthermore,

- (i) *If $t \mapsto t/m(t)$ is concave, then F is a constrained-efficient segmentation.*
- (ii) *If $t \mapsto t/m(t)$ is convex, then there exists a cutoff $\theta^c \in \Theta$, and posterior means $\underline{x} := \mathbb{E}_F[\theta | \theta \leq \theta^c]$ and $\bar{x} := \mathbb{E}_F[\theta | \theta \geq \theta^c]$ such that*

$$H(x) = \begin{cases} 0 & \text{if } x < \underline{x} \\ F(\theta^c) & \text{if } x \in [\underline{x}, \bar{x}) \\ 1 & \text{if } x \geq \bar{x} \end{cases}$$

is a constrained-efficient segmentation. In this case, $\tau_H^(\underline{x}) = 0 < \tau_H^*(\bar{x})$.*

¹¹Formally, $F(x) := \sigma_\mu^{PS}(\{\nu \in \Delta(\Theta) : \mathbb{E}_\nu[\theta] \leq x\})$ for all $x \in \Theta$.

¹²Note that the submarket tightness function is now defined over Θ , i.e., $\tau_H^* : \Theta \rightarrow \mathbb{R}_+$ maps a posterior mean $x \in \Theta$ to submarket tightness $\tau_H^*(x)$.

The above characterization has two implications for the constrained-efficient segmentation problem. First, it simplifies the non-linear optimization problem in (SB') by decomposing it into a linear programming problem and a fixed point problem involving the feasibility constraint. The linear programming problem can be addressed by applying tools from the information design literature (Dworczak and Martini, 2019; Kleiner et al., 2021; Kolotilin et al., 2022; Arieli et al., 2023). Specifically, the optimal segmentation can be characterized by a monotone partition, where the type space is divided into a countable number of connected intervals. Within each interval, the planner allocates each seller type to her own distinct submarket (perfect revelation); or allocates all seller types to same submarket (pooling); or randomly allocates all seller types to one of two submarkets (bi-pooling).

Second, under additional assumptions on $t/m(t)$, which represents the buyers' odds of meeting a seller (a 1-in- $t/m(t)$ chance), the constrained-efficient market segmentation takes one of two forms: perfect or binary market segmentation. To build intuition, recall that efficiency requires allocating more meetings to higher types. While the planner cannot directly control the number of meetings each seller type receives, it can indirectly influence how buyers distribute themselves across submarkets through the market segmentation. In particular, the planner can separate higher and lower seller types into distinct submarkets, which induces more buyers to self-select into the submarket for higher seller types, thereby increasing the number of meetings for those sellers

However, this strategy comes with trade-offs. While separating seller types ensures high-type sellers receive more meetings, the concavity of the meeting function implies a reduction in the total number of meetings in the search market. For this segmentation to enhance welfare, the surplus gained from increasing meetings for high-type sellers must outweigh the loss of surplus stemming from a lower number of overall meetings.

Proposition 5 demonstrates that such segmentation is welfare-improving when $t/m(t)$ is concave. Intuitively, as t increases, a buyer's odds of meeting a seller worsen, reflecting the congestion externality buyers face in tighter markets. When $t/m(t)$ is concave, this congestion externality worsens gradually as market tightness increases. Consequently, separating high and low seller types into distinct submarkets encourages many buyers to allocate themselves to the submarket for high-type sellers, mitigating the loss in welfare from a decrease in the total number of meet-

ings. Consequently, perfect segmentation emerges as constrained-efficient. Conversely, when $t/m(t)$ is convex, the congestion externality worsens more rapidly with increased submarket tightness, making such segmentation less attractive for the planner.

Let us conclude by noting that many meeting functions widely used in the search and matching literature satisfy the concavity or convexity conditions identified in [Proposition 5](#). For example, the mapping $t \mapsto t/m(t)$ is concave for CES meeting functions with parameter values $\rho \leq 1$, leading to perfect segmentation as the constrained-efficient outcome. In this case, buyers leverage granular ex-ante information about seller types to guide their choice of submarkets, thereby endogenously creating a directed search market where buyers learn a seller's type before meeting her.

In contrast, for Urn-ball meeting functions and for CES meeting functions with parameter values $\rho \geq 1$, the mapping $t \mapsto t/m(t)$ is convex. Similarly, while the generalized frictionless meeting function with $m(t) = \min\{\alpha, \beta t\}$ does not satisfy all the conditions of [Assumption 1](#), it inherits the properties of a CES meeting function, which approximates the frictionless meeting function as $\rho \rightarrow \infty$. As a result, the constrained-efficient segmentation in all three cases is a binary segmentation. In these cases, buyers have minimal ex-ante information they can use to direct their search. Instead, a random search market emerges, with buyers learning a seller's type only after meeting her.

A. Appendix

A.1. Proofs

Proof of Proposition 1. Let us fix a market segmentation σ . Let $\mathcal{L}^1(\Delta(\Theta), \sigma_\mu)$ be the space of σ_μ -integrable functions from $\Delta(\Theta)$ to \mathbb{R} . Consider the following relaxed problem:

$$\begin{aligned} \max_{\substack{\tau \in \mathcal{L}^1(\Delta(\Theta), \sigma_\mu) \\ \tau \geq 0}} & k \int_{\Delta(\Theta)} m(\tau(\nu)) \mathbb{E}_\nu[\theta] \sigma_\mu(d\nu) \\ \text{s.t.} & k \int_{\Delta(\Theta)} \tau(\nu) \sigma_\mu(d\nu) \leq 1. \end{aligned}$$

Notice that the constraint in the relaxed problem must bind at the optimum since $m(t)$ is a strictly increasing function. Hence, any solution to the relaxed problem is also a solution to (FB- σ).

Notice that both the objective and the constraint in the relaxed problem are concave functionals. Thus, from Luenberger (1997) (Section 8.3, Theorem 1; Section 8.4, Theorem 2), τ_σ^{FB} is a solution to (FB- σ) if and only if there exists a Lagrange multiplier $\eta_\sigma > 0$ such that τ_σ^{FB} solves:

$$\max_{\tau \in \mathcal{L}^1(\Delta(\Theta), \sigma_\mu)} L(\tau, \eta_\sigma) := k \int_{\Delta(\Theta)} \left[m(\tau(\nu)) \mathbb{E}_\nu[\theta] - \eta_\sigma \tau(\nu) \right] \sigma_\mu(d\nu) + \eta_\sigma.$$

We can maximize the Lagrangian by pointwise maximizing the integrand, which depends on σ only indirectly through η_σ . Hence, we can first pointwise maximize $L(\tau, \eta)$ for an arbitrary $\eta > 0$, and then find the value η_σ that makes the constraint in (FB- σ) bind.

To that end, given $\eta > 0$, let us find a function $\psi : \Delta(\Theta) \times (0, \infty) \rightarrow \mathbb{R}_+$ such that $m'(\psi(\nu, \eta)) \mathbb{E}_\nu[\theta] \leq \eta$ for all $\nu \in \text{supp}(\sigma_\mu)$ with equality if $\psi(\nu, \eta) > 0$. These first-order conditions are also sufficient as $m(t) \mathbb{E}_\nu[\theta] - \eta t$ is concave in t for all $\nu \in \Delta(\Theta)$. Recall that $m'(t)$ is strictly decreasing in t with $m'(0) = \beta$. Thus, set $\psi(\nu, \eta) = 0$ for all $\nu \in \text{supp}(\sigma_\mu)$ with $\beta \mathbb{E}_\nu[\theta] \leq \eta$. Furthermore, for all $\nu \in \text{supp}(\sigma_\mu)$ with $\beta \mathbb{E}_\nu[\theta] > \eta$, set $m'(\psi(\nu, \eta)) \mathbb{E}_\nu[\theta] = \eta$. As a result, recalling the definition of $f : [1/\beta, \infty) \rightarrow \mathbb{R}_+$,

we have

$$\psi(\nu, \eta) := \begin{cases} 0 & \text{if } \beta \mathbb{E}_\nu[\theta] \leq \eta \\ f\left(\frac{\mathbb{E}_\nu[\theta]}{\eta}\right) & \text{if } \beta \mathbb{E}_\nu[\theta] > \eta \end{cases}.$$

Since f is a continuous and strictly increasing function with $\lim_{y \rightarrow \infty} f(y) = \infty$ and $f(1/\beta) = 0$, we have:

- (i) the mapping $(\nu, \eta) \mapsto \psi(\nu, \eta)$ is continuous in both arguments,
- (ii) for each $\eta > 0$, $\psi(\nu, \eta) \geq \psi(\nu', \eta)$ whenever $\mathbb{E}_\nu[\theta] \geq \mathbb{E}_{\nu'}[\theta]$,
- (iii) for each $\nu \in \text{supp}(\sigma_\mu)$, $\psi(\nu, \eta) \leq \psi(\nu, \eta')$ whenever $\eta > \eta' > 0$, and
- (iv) for σ_μ -almost all ν , $\lim_{\eta \rightarrow 0} \psi(\nu, \eta) = \infty$, and $\psi(\nu, \eta) = 0$ for all $\eta \geq \beta$.

Consequently, $\psi(\nu, \eta) \leq \psi(\delta_1, \eta) < \infty$ for all $\nu \in \text{supp}(\sigma_\mu)$ and all $\eta > 0$, which in turn implies $\psi(\cdot, \eta) \in \mathcal{L}^1(\Delta(\Theta), \sigma_\mu)$ for all $\eta > 0$.

Next, let us show that there exists a unique multiplier $\eta_\sigma \in (0, \beta)$ such that η_σ along with $\psi(\cdot, \eta_\sigma)$ lead to a binding feasibility constraint in (FB- σ). To that end, define $\Psi : (0, \infty) \rightarrow \mathbb{R}_+$ to be the mapping given by

$$\Psi(\eta) := k \int_{\Delta(\Theta)} \psi(\nu, \eta) \sigma_\mu(d\nu).$$

Clearly, the feasibility constraint in (FB- σ) binds at η_σ if $\Psi(\eta_\sigma) = 1$.

Observe that Ψ is a continuous and weakly decreasing function because $\psi(\nu, \eta)$ is continuous and weakly decreasing in η for each $\nu \in \text{supp}(\sigma_\mu)$. Moreover, $\Psi(\eta) = 0$ for all $\eta \geq \beta$, and from the monotone convergence theorem,

$$\lim_{\eta \rightarrow 0} \Psi(\eta) = k \int_{\Delta(\Theta)} \lim_{\eta \rightarrow 0} \psi(\nu, \eta) \sigma_\mu(d\nu) = \infty.$$

Thus, there exists a unique $\eta_\sigma \in (0, \beta)$ such that $\Psi(\eta_\sigma) = 1$. The optimal submarket tightness is defined by $\tau_\sigma^{FB} := \psi(\cdot, \eta_\sigma)$. Furthermore, any τ that solves (FB- σ) must have $\tau = \tau_\sigma^{FB}$ for σ_μ -almost everywhere. In other words, the optimal submarket tightness for a given segmentation σ is unique up to the σ_μ -null sets. \square

Proof of Proposition 2. Fix any market segmentation σ , and let τ_σ^{FB} be the solution to (FB- σ). The proof proceeds by showing that, given the pair $(\sigma, \tau_\sigma^{FB})$, there exists another pair $(\sigma^{PS}, \tilde{\tau})$ that achieves a (weakly) higher surplus for each type- θ seller. To that end, define a new submarket tightness function $\tilde{\tau} : \Delta(\Theta) \rightarrow \mathbb{R}_+$ given by

$$\tilde{\tau}(\nu) = \begin{cases} \int_{\Delta(\Theta)} \tau_\sigma^{FB}(\nu') \sigma(d\nu', \theta) & \text{if } \nu = \delta_\theta \\ \tau_\sigma^{FB}(\nu) & \text{otherwise} \end{cases},$$

and notice that

$$\int_{\Delta(\Theta)} \tilde{\tau}(\nu) \sigma_\mu^{PS}(d\nu) = \int_{\Theta} \tilde{\tau}(\delta_\theta) \mu(d\theta) = \int_{\Delta(\Theta)} \tau_\sigma^{FB}(\nu) \sigma_\mu(d\nu),$$

where the first equality follows from the definition of σ^{PS} and the second equality follows by construction of $\tilde{\tau}$. Therefore, $\tilde{\tau} \in \mathcal{L}^1(\Delta(\Theta), \sigma_\mu^{PS})$ and the pair $(\sigma^{PS}, \tilde{\tau})$ satisfies (2).

Comparing the expected surplus generated from a type- θ seller under $(\sigma, \tau_\sigma^{FB})$ to that under $(\sigma^{PS}, \tilde{\tau})$, we have

$$\begin{aligned} k \cdot \theta \int_{\Delta(\Theta)} m(\tau_\sigma^{FB}(\nu)) \sigma(d\nu, \theta) &\leq k \cdot \theta \cdot m \left(\int_{\Delta(\Theta)} \tau_\sigma^{FB}(\nu) \sigma(d\nu, \theta) \right) \\ &= k \cdot \theta \cdot m(\tilde{\tau}(\delta_\theta)) \\ &= k \cdot \theta \int_{\Delta(\Theta)} m(\tilde{\tau}(\nu)) \sigma^{PS}(d\nu, \theta), \end{aligned}$$

where the inequality follows from the concavity of m , the first equality follows by definition of $\tilde{\tau}$, and the last follows by the definition of σ^{PS} .

Thus, for any arbitrary market segmentation σ , the pair $(\sigma, \tau_\sigma^{FB})$ generates a (weakly) lower expected surplus than the pair $(\sigma^{PS}, \tilde{\tau})$. Of course, $(\sigma^{PS}, \tilde{\tau})$ generates a (weakly) lower expected surplus than the pair $(\sigma^{PS}, \tau_{\sigma^{PS}}^{FB})$, because the latter solves (FB- σ) when $\sigma = \sigma^{PS}$. As a result, the pair $(\sigma^{PS}, \tau_{\sigma^{PS}}^{FB})$ is a solution to (FB). \square

Proof of Proposition 3. Let us first establish a sufficient and necessary condition for a submarket to be active. We shall then use this result to show that an essentially-unique equilibrium exists for any given market segmentation.

Lemma 1. *Suppose $(\tau_\sigma^*, u_\sigma^*)$ is a search equilibrium for a market segmentation σ . For σ_μ -almost all ν , $\tau_\sigma^*(\nu) > 0$ if and only if $\beta \mathbb{E}_\nu[\lambda(\theta)\theta] > u_\sigma^*$.*

Proof of Lemma 1.

(*Only-if direction*): Suppose submarket $\nu \in \text{supp}(\sigma_\mu)$ is active, i.e., $\tau_\sigma^*(\nu) > 0$, which implies

$$u_\sigma^* = \frac{m(\tau_\sigma^*(\nu))}{\tau_\sigma^*(\nu)} \mathbb{E}_\nu[\lambda(\theta)\theta] < \beta \mathbb{E}_\nu[\lambda(\theta)\theta],$$

where the equality follows from (3), and the inequality follows from the fact β is the highest probability with which a buyer meets a seller.

(*If direction*): Suppose submarket $\nu \in \text{supp}(\sigma_\mu)$ is inactive, i.e., $\tau_\sigma^*(\nu) = 0$, which implies

$$u_\sigma^* \geq \frac{m(0)}{0} \mathbb{E}_\nu[\lambda(\theta)\theta] = \beta \mathbb{E}_\nu[\lambda(\theta)\theta],$$

where the inequality follows from (3), and the equality follows by definition that $m(0)/0 = \beta$. Thus, from the contrapositive, $\beta \mathbb{E}_\nu[\lambda(\theta)\theta] > u_\sigma^*$ implies $\tau_\sigma^*(\nu) > 0$. \square

Let us now turn to proving the existence of an essentially-unique equilibrium. By Lemma 1, any search equilibrium $(\tau_\sigma^*, u_\sigma^*)$ satisfies $\tau_\sigma^*(\nu) = 0$ for σ_μ -almost all ν such that $\beta \mathbb{E}_\nu[\lambda(\theta)\theta] \leq u_\sigma^*$, and $\tau_\sigma^*(\nu) > 0$ for σ_μ -almost all ν such that $\beta \mathbb{E}_\nu[\lambda(\theta)\theta] > u_\sigma^*$. In the latter case, (3) implies that

$$\frac{\tau_\sigma^*(\nu)}{m(\tau_\sigma^*(\nu))} = \frac{\mathbb{E}_\nu[\lambda(\theta)\theta]}{u_\sigma^*},$$

which from the definition of g implies that $\tau_\sigma^*(\nu) = g(\mathbb{E}_\nu[\lambda(\theta)\theta]/u_\sigma^*)$. In other words, there exists a mapping $\phi : \Delta(\Theta) \times (0, \infty) \rightarrow \mathbb{R}_+$ given by

$$\phi(\nu, u) := \begin{cases} 0 & \text{if } \beta \mathbb{E}_\nu[\lambda(\theta)\theta] \leq u \\ g\left(\frac{\mathbb{E}_\nu[\lambda(\theta)\theta]}{u}\right) & \text{if } \beta \mathbb{E}_\nu[\lambda(\theta)\theta] > u \end{cases}$$

such that $\tau_\sigma^* = \phi(\cdot, u_\sigma^*)$ for σ_μ -almost everywhere.

Since g is a continuous and strictly increasing function with $\lim_{y \rightarrow \infty} g(y) = \infty$ and $g(1/\beta) = 0$, we have:

- (i) the mapping $(\nu, u) \mapsto \phi(\nu, u)$ is continuous in both arguments,

- (ii) for each $u > 0$, $\phi(\nu, u) \geq \psi(\nu', u)$ whenever $\mathbb{E}_\nu[\lambda(\theta)\theta] \geq \mathbb{E}_{\nu'}[\lambda(\theta)\theta]$, and
- (iii) for each $\nu \in \text{supp}(\sigma_\mu)$, $\phi(\nu, u) \leq \phi(\nu, u')$ whenever $u > u' > 0$, and
- (iv) for σ_μ -almost all ν , $\lim_{u \rightarrow 0} \phi(\nu, u) = \infty$, and $\psi(\nu, u) = 0$ for all $u \geq \beta \sup_{\theta \in \Theta} \lambda(\theta)\theta$.

Consequently, $\phi(\nu, u) \leq \sup_{\nu' \in \Delta(\Theta)} \phi(\nu', u) \leq g(1/u) < \infty$ for all $\nu \in \text{supp}(\sigma_\mu)$, which in turn implies $\phi(\cdot, u) \in \mathcal{L}^1(\Delta(\Theta), \sigma_\mu)$ for all $u > 0$.

Next, observe that a pair $(\tau_\sigma^*, u_\sigma^*)$ jointly satisfies (3) and (4) if and only if u_σ^* is a fixed point of the function $\Phi : (0, \infty) \rightarrow \mathbb{R}_+$ given by

$$\Phi(u) := k \int_{\Delta(\Theta)} m(\phi(\nu, u)) \mathbb{E}_\nu[\lambda(\theta)\theta] \sigma_\mu(d\nu).$$

The function $\Phi(u)$ is a continuous and weakly decreasing function because $\phi(\nu, u)$ is continuous and weakly decreasing in u for each $\nu \in \text{supp}(\sigma_\mu)$. Moreover, $\Phi(u) = 0$ for all $u \geq \beta \sup_{\theta \in \Theta} \lambda(\theta)\theta$, and

$$\begin{aligned} \lim_{u \rightarrow 0} \Phi(u) &= k \int_{\Delta(\Theta)} \lim_{u \rightarrow 0} m(\phi(\nu, u)) \mathbb{E}_\nu[\lambda(\theta)\theta] \sigma_\mu(d\nu) \\ &= \alpha k \mathbb{E}_\mu[\lambda(\theta)\theta] \\ &> 0, \end{aligned}$$

where the first equality follows from the dominated convergence theorem, the second equality follows because $\lim_{u \rightarrow 0} \phi(\nu, u) = \infty$ for all ν such that $\mathbb{E}_\nu[\lambda(\theta)\theta] > 0$, and because $\lim_{t \rightarrow \infty} m(t) = \alpha$, and the inequality follows from [Assumption 2](#).

Thus far, we have shown that $\Phi(u)$ is a continuous and (weakly) decreasing function with $\lim_{u \rightarrow \beta \sup_{\theta \in \Theta} \lambda(\theta)\theta} \Phi(u) = 0 < \lim_{u \rightarrow 0} \Phi(u)$. Therefore, it has a unique fixed point in the interval $(0, \beta \sup_{\theta \in \Theta} \lambda(\theta)\theta)$. We have thus established the existence of an essentially-unique pair $(\tau_\sigma^*, u_\sigma^*)$ with $u_\sigma^* = \Phi(u_\sigma^*)$ and $\tau_\sigma^*(\cdot) = \phi(\cdot, u_\sigma^*)$, such that $(\tau_\sigma^*, u_\sigma^*)$ jointly satisfies (3) and (4). In other words, there exists an essentially-unique search equilibrium $(\tau_\sigma^*, u_\sigma^*)$ for each market segmentation σ . \square

Proof of [Proposition 4](#). (Only-if direction): Suppose the pair $(\sigma^{PS}, \tau_{\sigma^{PS}}^*)$ solves (FB). By [Proposition 1](#), $\tau_{\sigma^{PS}}^* = \tau_{\sigma^{PS}}^{FB}$ for σ_μ^{PS} -almost everywhere, or equivalently, $\tau_{\sigma^{PS}}^*(\delta_\theta) = \tau_{\sigma^{PS}}^{FB}(\delta_\theta)$ for μ -almost all θ .

From [Proposition 1](#), $\tau_{\sigma^{PS}}^{FB}(\delta_\theta) = 0$ for all $\theta \leq \eta_{\sigma^{PS}}/\beta$, and $\tau_{\sigma^{PS}}^{FB}(\delta_\theta) > 0$ for all $\theta > \eta_{\sigma^{PS}}/\beta$. Similarly, from [Proposition 3](#), $\tau_{\sigma^{PS}}^*(\delta_\theta) = 0$ for all θ such that $\lambda(\theta)\theta \leq u_{\sigma^{PS}}^*/\beta$, and $\tau_{\sigma^{PS}}^*(\delta_\theta) > 0$ for all θ such that $\lambda(\theta)\theta > u_{\sigma^{PS}}^*/\beta$. Since the set of active and inactive submarkets coincide almost everywhere under the two submarket tightness functions, we have that $\theta \leq \eta_{\sigma^{PS}}/\beta$ if and only if $\lambda(\theta)\theta \leq u_{\sigma^{PS}}^*/\beta$. By the continuity of the surplus-splitting function, we can conclude that

$$u_{\sigma^{PS}}^* = \lambda\left(\frac{\eta_{\sigma^{PS}}}{\beta}\right)\eta_{\sigma^{PS}}.$$

Thus, for all $\theta \leq \eta_{\sigma^{PS}}/\beta$,

$$\lambda(\theta)\theta \leq \lambda\left(\frac{\eta_{\sigma^{PS}}}{\beta}\right)\frac{\eta_{\sigma^{PS}}}{\beta}.$$

Moreover, for all $\theta > \eta_{\sigma^{PS}}/\beta$, [Proposition 1](#) implies that

$$m'(\tau_{\sigma^{PS}}^{FB}(\delta_\theta))\theta = \eta_{\sigma^{PS}},$$

while [Proposition 3](#) implies that

$$\frac{m(\tau_{\sigma^{PS}}^*(\delta_\theta))}{\tau_{\sigma^{PS}}^*(\delta_\theta)}\lambda(\theta)\theta = \lambda\left(\frac{\eta_{\sigma^{PS}}}{\beta}\right)\eta_{\sigma^{PS}}.$$

Combining the above two expressions and noting that $\tau_{\sigma^{PS}}^{FB} = \tau_{\sigma^{PS}}^*$ almost everywhere yields

$$\lambda(\theta) = \lambda\left(\frac{\eta_{\sigma^{PS}}}{\beta}\right) \cdot \varepsilon\left(f\left(\frac{\theta}{\eta_{\sigma^{PS}}}\right)\right)$$

for all $\theta > \eta_{\sigma^{PS}}/\beta$.

(If direction): Suppose the surplus-splitting function, λ , satisfies the properties given in the proposition. We first show that $(\tau_{\sigma^{PS}}^{FB}, \lambda(\eta_{\sigma^{PS}}/\beta) \cdot \eta_{\sigma^{PS}})$ is a search equilibrium of σ^{PS} .

To that end, for $\theta \leq \eta_{\sigma^{PS}}/\beta$, we have

$$U(\delta_\theta; \tau_{\sigma^{PS}}^{FB}, \sigma^{PS}) = \underbrace{\frac{m(0)}{0}}_{=\beta} \lambda(\theta)\theta \leq \lambda\left(\frac{\eta_{\sigma^{PS}}}{\beta}\right)\eta_{\sigma^{PS}},$$

where the equality follows from the fact that $\tau_{\sigma^{PS}}^{FB}(\delta_\theta) = 0$ for $\theta \leq \eta_{\sigma^{PS}}/\beta$ ([Proposi-](#)

tion 1), and the inequality follows by the assumption on λ given in the proposition.

For $\theta > \eta_{\sigma^{PS}}/\beta$, we have

$$\begin{aligned} U(\delta_\theta; \tau_{\sigma^{PS}}^{FB}, \sigma^{PS}) &= \frac{m\left(f\left(\frac{\theta}{\eta_{\sigma^{PS}}}\right)\right)}{f\left(\frac{\theta}{\eta_{\sigma^{PS}}}\right)} \lambda(\theta)\theta \\ &= m'\left(f\left(\frac{\theta}{\eta_{\sigma^{PS}}}\right)\right) \lambda\left(\frac{\eta_{\sigma^{PS}}}{\beta}\right) \theta \\ &= \lambda\left(\frac{\eta_{\sigma^{PS}}}{\beta}\right) \eta_{\sigma^{PS}}, \end{aligned}$$

where the first equality follows from the characterization of $\tau_{\sigma^{PS}}^{FB}(\delta_\theta)$ for $\theta > \eta_{\sigma^{PS}}/\beta$ (Proposition 1), the second equality follows by the assumption on λ given in the proposition, and the last equality follows from the fact that f is the inverse of the mapping $t \mapsto 1/m'(t)$. Therefore, the pair $(\tau_{\sigma^{PS}}^{FB}, \lambda(\eta_{\sigma^{PS}})\eta_{\sigma^{PS}})$ satisfies (3) of Definition 1.

Additionally,

$$\begin{aligned} k \int_{\Delta(\Theta)} m(\tau_{\sigma^{PS}}^{FB}(\delta_\theta)) \mathbb{E}_\nu[\lambda(\theta)\theta] \sigma_\mu^{PS}(d\nu) &= k \int_{\eta_{\sigma^{PS}}/\beta}^1 m\left(f\left(\frac{\theta}{\eta_{\sigma^{PS}}}\right)\right) \lambda(\theta)\theta \mu(d\theta) \\ &= \lambda\left(\frac{\eta_{\sigma^{PS}}}{\beta}\right) \eta_{\sigma^{PS}} \cdot k \int_{\eta_{\sigma^{PS}}/\beta}^1 f\left(\frac{\eta_{\sigma^{PS}}}{\theta}\right) \mu(d\theta) \\ &= \lambda\left(\frac{\eta_{\sigma^{PS}}}{\beta}\right) \eta_{\sigma^{PS}}, \end{aligned}$$

where the first equality follows from the definition of σ^{PS} and the characterization of $\tau_{\sigma^{PS}}^{FB}$ in Proposition 1, the second equality follows by the assumption on λ given in Proposition 4 and the fact that f is the inverse of the mapping $t \mapsto 1/m'(t)$, and the last equality follows by the construction of $\eta_{\sigma^{PS}}$ as given in Proposition 1. Thus, the pair $(\tau_{\sigma^{PS}}^{FB}, \lambda(\eta_{\sigma^{PS}}/\beta) \cdot \eta_{\sigma^{PS}})$ also satisfies (4) of Definition 1.

We can therefore conclude that $(\tau_{\sigma^{PS}}^{FB}, \lambda(\eta_{\sigma^{PS}}/\beta) \cdot \eta_{\sigma^{PS}})$ is a search equilibrium of σ^{PS} . However, $(\tau_{\sigma^{PS}}^*, u_{\sigma^{PS}}^*)$ is the essentially-unique search equilibrium of σ^{PS} . Thus, $\tau_{\sigma^{PS}}^* = \tau_{\sigma^{PS}}^{FB}$ for σ_μ^{PS} -almost everywhere, and $u_{\sigma^{PS}}^* = \lambda(\eta_{\sigma^{PS}}/\beta) \cdot \eta_{\sigma^{PS}}$. Since $(\sigma^{PS}, \tau_{\sigma^{PS}}^{FB})$ solves (FB), the almost-everywhere equivalence between $\tau_{\sigma^{PS}}^*$ and $\tau_{\sigma^{PS}}^{FB}$ implies that

$(\sigma^{PS}, \tau_{\sigma^{PS}}^*)$ also solves (FB), giving us the desired result. \square

Proof of Proposition 5. Let us first establish some facts that will be useful in the proof: Under Assumption 3, $\beta \sup_{\theta \in \Theta} \lambda(\theta)\theta = \beta\ell$. With some abuse of notation, for $x \in \Theta$ and $u > 0$, let

$$\phi(x, u) := \begin{cases} 0 & \text{if } \beta\ell x \leq u \\ g\left(\frac{\ell x}{u}\right) & \text{if } \beta\ell x > u \end{cases}$$

which is simply a reformulation of the function $\phi(\nu, u)$ as defined in the proof of Proposition 3 into the current environment where payoffs depend only on the posterior mean.

Given $u > 0$, consider the following linear program:

$$\max_{\hat{H} \in \text{MPC}(F)} k \int_{\Theta} \phi(x, u) d\hat{H}(x). \quad (\text{LP-}u)$$

Let $V(u)$ be the value function corresponding to (LP- u), and let $\mathcal{H}(u)$ be the set of maximizers. Since the objective is a continuous function and $\text{MPC}(F)$ is compact, $V(u)$ is continuous in u , and $\mathcal{H}(u)$ is non-empty and compact-valued for each $u > 0$. Importantly, by construction, H satisfies (5) if and only if $H \in \mathcal{H}(u_H^*)$.

Observe that V is a (weakly) decreasing function as $\phi(x, u)$ is a (weakly) decreasing function of u for all $x \in \Theta$. Furthermore, for almost all $x \in \Theta$, $\phi(x, u)$ grows arbitrarily large as $u \rightarrow 0$ and $\phi(x, u) = 0$ for all $u \geq \beta\ell$, which implies that $\lim_{u \rightarrow 0} V(u) = \infty$ and $V(u) = 0$ for all $u \geq \beta\ell$. Thus, there is a unique point $\bar{u} \in (0, \beta\ell)$ such that $V(\bar{u}) = 1$.

Lemma 2. *Consider a market segmentation $H \in \text{MPC}(F)$, and let (τ_H^*, u_H^*) be its essentially-unique search equilibrium. Then $u_H^* \leq \bar{u}$. Furthermore, $H \in \mathcal{H}(u_H^*)$ if and only if $u_H^* = \bar{u}$.*

Proof of Lemma 2. For each $H \in \text{MPC}(F)$,

$$V(\bar{u}) = 1 = k \int_{\Theta} \phi(x, u_H^*) dH(x) \leq V(u_H^*),$$

where the first equality follows by construction of \bar{u} , the second equality follows from (2), and the inequality follows because V is the value function of (LP- u). Since V is a weakly decreasing function, we have $\bar{u} \geq u_H^*$, establishing the first desired result.

Furthermore, by construction, the above inequality holds with equality if and only if $H \in \mathcal{H}(u_H^*)$. However, since \bar{u} is the unique value for which $V(u) = 1$, the above inequality holds with equality if and only if $u_H^* = \bar{u}$. This establishes the second desired result. \square

Let us now prove the general characterization for constrained-efficient market segmentations given in (5), and then prove the characterization for the special cases when $t/m(t)$ is convex or concave.

General case:

(*Only-if* direction): Suppose H is a constrained-efficient market segmentation, i.e., H solves (SB'). We want to show that $H \in \mathcal{H}(u_H^*)$. By Lemma 2, this is equivalent to showing that $u_H^* = \bar{u}$. Furthermore, Lemma 2 already establishes that $\bar{u} \geq u_H^*$. Thus, to show the desired “only if” direction, it suffices to show that $u_H^* \geq \bar{u}$.

Notice that under Assumption 3, the buyers’ payoff is equivalent to expected total surplus scaled down by a constant ℓ . Therefore, because H solves (SB'), we have $u_H^* \geq u_{\hat{H}}^*$ for all $\hat{H} \in \text{MPC}(F)$. In particular—recalling that $\mathcal{H}(\bar{u})$ is non-empty—for each $\hat{H} \in \mathcal{H}(\bar{u})$, we have $u_H^* \geq u_{\hat{H}}^* = \bar{u}$, where the last equality follows from Lemma 2.

(*If* direction): Suppose H satisfies (5), i.e., $H \in \mathcal{H}(u_H^*)$. By Lemma 2, $\bar{u} = u_H^* \geq u_{\hat{H}}^*$ for all $\hat{H} \in \text{MPC}(F)$. In other words, for all $\hat{H} \in \text{MPC}(F)$,

$$\underbrace{\ell \cdot k \int_{\Theta} m(\tau_H^*(x)) x dH(x)}_{=u_H^*} \geq \underbrace{\ell \cdot k \int_{\Theta} m(\tau_{\hat{H}}^*(x)) x d\hat{H}(x)}_{=u_{\hat{H}}^*},$$

where the equalities follow by (4). The above inequality then establishes the desired result: H solves (SB').

We have thus far shown that H is a constrained-efficient market segmentation if and only if $H \in \mathcal{H}(\bar{u}) \equiv \mathcal{H}(u_H^*)$. This concludes the characterization for the general statement.

Let us proceed by showing that when the mapping $t \mapsto t/m(t)$ is concave, then

$F \in \mathcal{H}(u)$ for all $u \in (0, \beta\ell)$. Similarly, when the mapping $t \mapsto t/m(t)$ is convex, we show that for each $u \in (0, \beta\ell)$, there exists a binary market segmentation in $\mathcal{H}(u)$.

First, suppose the mapping $t \mapsto t/m(t)$ is **concave**, and therefore, g is convex. In this case, $\phi(x, u) = g(\ell \cdot x/u) \mathbb{1}_{[\ell\beta x > u]}$ is convex in x for all $u \in (0, \beta\ell)$.¹³ As a result, for all $H', H'' \in \text{MPC}(F)$ where H' is a mean-preserving contraction of H'' , we have

$$k \int_{\Theta} \phi(x, u) dH'(x) \leq k \int_{\Theta} \phi(x, u) dH''(x).$$

Therefore, $F \in \mathcal{H}(u)$ for all $u \in (0, \beta\ell)$. Consequently, $F \in \mathcal{H}(\bar{u})$, i.e., there exists a perfect market segmentation is constrained-efficient, establishing Point (i) of [Proposition 5](#).

Second, suppose the mapping $t \mapsto t/m(t)$ is **convex**, and therefore, g is concave. In this case, $\phi(x, u) = g(\ell \cdot x/u) \mathbb{1}_{[\ell\beta x > u]}$ is piece-wise concave in x for all $u \in (0, \beta\ell)$. Specifically, it is constant over $[0, u/(\beta\ell)]$ and concave over $[u/(\beta\ell), 1]$.¹⁴

Define the functions $\underline{X} : \Theta \rightarrow \Theta$ and $\bar{X} : \Theta \rightarrow \Theta$ given by $\underline{X}(\theta) := \mathbb{E}_F[\tilde{\theta} | \tilde{\theta} \leq \theta]$ and $\bar{X}(\theta) := \mathbb{E}_F[\tilde{\theta} | \tilde{\theta} \geq \theta]$, respectively. As F is assumed to be absolutely continuous, both \underline{X} and \bar{X} are continuous and strictly increasing functions. Furthermore, it can be readily checked that the mapping $\theta \mapsto \theta - \underline{X}(\theta)$ is positive and strictly increasing over $(0, 1)$, while the mapping $\theta \mapsto \theta - \bar{X}(\theta)$ is negative and strictly increasing over $(0, 1)$.

For each $u \in (0, \beta\ell)$, let $\underline{\theta}_u := \min\{\theta \in \Theta : \bar{X}(\theta) \geq u/(\beta\ell)\}$, which is well-defined because \bar{X} is continuous and strictly increasing, and $\{\theta \in \Theta : \bar{X}(\theta) \geq u/(\beta\ell)\}$ is non-empty.¹⁵ Define $G_u : [\underline{\theta}_u, 1] \rightarrow \mathbb{R}$ as the function given by

$$G_u(\theta) := g\left(\frac{\ell \cdot \bar{X}(\theta)}{u}\right) + g'\left(\frac{\ell \cdot \bar{X}(\theta)}{u}\right) \frac{\ell}{u} (\theta - \bar{X}(\theta)),$$

which is a continuous and strictly increasing function because (a) both g and \bar{X} are

¹³Recall that $\lim_{y \rightarrow 1/\beta} g(y) = 0$.

¹⁴[Kolotilin et al. \(2022\)](#) describe such a ϕ function as ‘‘S-shaped.’’ They show that if ϕ is continuously differentiable and S-shaped, there exists a solution to (LP- u) that is an upper-censorship: perfectly segment types below a cutoff and pool all the types above the cutoff. However, their results cannot be directly applied here as the $\phi(x, u)$ is not differentiable at $x = \frac{u}{\beta\ell}$. Additionally, [Proposition 5](#) goes further than an upper censorship; it claims that a binary segmentation is sufficient to obtain constrained efficiency.

¹⁵For example, $\bar{X}(1) = 1 > u/(\beta\ell)$.

continuous and strictly increasing functions, (b) g' is a (weakly) decreasing function by concavity of g , and (c) the mapping $\theta \mapsto \theta - \bar{X}(\theta)$ is negative and strictly increasing over $(0, 1)$.

Next, define the cutoff $\theta_u^c := \min\{\theta \in [\underline{\theta}_u, 1] : G_u(\theta) \geq 0\}$, which is well-defined because G_u is continuous and strictly increasing, and $\{\theta \in [\underline{\theta}_u, 1] : G_u(\theta) \geq 0\}$ is non-empty.¹⁶

Lemma 3. *For each $u \in (0, \beta\ell)$,*

$$\underline{X}(\theta_u^c) \leq \frac{u}{\beta\ell} < \bar{X}(\theta_u^c).$$

Proof of Lemma 3. Fix an arbitrary $u \in (0, \beta\ell)$. First, notice that $\bar{X}(u/(\beta\ell)) > u/(\beta\ell)$ by construction. Since $\underline{\theta}_u$ is the smallest type in Θ that satisfies $\bar{X}(\theta) \geq u/(\beta\ell)$, we have $u/(\beta\ell) \geq \underline{\theta}_u$.

Next, notice that

$$\begin{aligned} G\left(\frac{u}{\beta\ell}\right) &= g\left(\frac{\ell \cdot \bar{X}\left(\frac{u}{\beta\ell}\right)}{u}\right) + g'\left(\frac{\ell \cdot \bar{X}\left(\frac{u}{\beta\ell}\right)}{u}\right) \frac{\ell}{u} \left(\frac{u}{\beta\ell} - \bar{X}\left(\frac{u}{\beta\ell}\right)\right) \\ &\geq g\left(\frac{\ell \cdot \frac{u}{\beta\ell}}{u}\right) \\ &= 0, \end{aligned}$$

where the inequality follows from the concavity of g , and the last equality follows because $g(1/\beta) = 0$. Since θ_u^c is the smallest type in $[\underline{\theta}_u, 1]$ that satisfies $G_u(\theta) \geq 0$, we have $u/(\beta\ell) \geq \theta_u^c$.

By construction, $\underline{X}(\theta) \leq \theta$ for all $\theta \in \Theta$. Thus, $\underline{X}(\theta_u^c) \leq \theta_u^c \leq u/(\beta\ell)$, which establishes the first inequality in the lemma.

Also by construction, $\theta_u^c \geq \underline{\theta}_u$, which implies that $\bar{X}(\theta_u^c) \geq \bar{X}(\underline{\theta}_u) \geq u/(\beta\ell)$. Suppose, for the sake of contradiction, $\bar{X}(\theta_u^c) = u/(\beta\ell)$. Then

$$G_u(\theta_u^c) = \underbrace{g\left(\frac{\ell \cdot \frac{u}{\beta\ell}}{u}\right)}_{=0} + \underbrace{g'\left(\frac{\ell \cdot \frac{u}{\beta\ell}}{u}\right)}_{>0} \underbrace{\frac{\ell}{u} \left(\theta_u^c - \bar{X}(\theta_u^c)\right)}_{<0} < 0,$$

¹⁶For example, $1 - \bar{X}(1) = 0$, so $G_u(1) = g(\ell/u) > 0$ for each $u \in (0, \beta\ell)$. The inequality follows because g is a strictly increasing function with $g(1/\beta) = 0$.

which contradicts the fact that, by definition, $G_u(\theta_u^c) \geq 0$. To avoid the contradiction, we must have $\bar{X}(\theta_u^c) > u/(\beta\ell)$, establishing the second inequality in the lemma. \square

Finally, consider the binary posterior-mean distribution $H_u \in \text{MPC}(F)$ given by

$$H_u(x) = \begin{cases} 0 & \text{if } x < \underline{X}(\theta_u^c) \\ F(\theta^c) & \text{if } x \in [\underline{X}(\theta_u^c), \bar{X}(\theta_u^c)] \\ 1 & \text{if } x \geq \bar{X}(\theta_u^c) \end{cases} .$$

From Theorem 1 of [Dworczak and Martini \(2019\)](#) (DM henceforth), H_u solves (LP- u) if there exists a convex function $p_u : \Theta \rightarrow \mathbb{R}$ such that

- (a) $\phi(\cdot, u) \leq p_u$ pointwise,
- (b) $\text{supp}(H_u) \subseteq \{x \in \Theta : \phi(x, u) = p_u(x)\}$, and
- (c) $\mathbb{E}_{H_u}[p_u(x)] = \mathbb{E}_F[p_u(x)]$.

Consider the function $p : \Theta \rightarrow \mathbb{R}$ given by

$$p_u(x) = \begin{cases} 0 & \text{if } x < \theta_u^c \\ g\left(\frac{\ell \cdot \bar{X}(\theta_u^c)}{u}\right) + g'\left(\frac{\ell \cdot \bar{X}(\theta_u^c)}{u}\right) \frac{\ell}{u} (x - \bar{X}(\theta_u^c)) & \text{if } x \geq \theta_u^c \end{cases} .$$

Notice that $p_u(x)$ is (weakly) increasing and convex as it is the upper envelope of two affine functions. Clearly, $p_u(x) = \phi(x, u)$ for all $x < \theta_u^c$. For $x \in [\theta_u^c, u/(\beta\ell)]$, we have

$$\phi(x, u) = 0 \leq G_u(\theta_u^c) = p_u(\theta_u^c) \leq p_u(x),$$

where the first inequality follows by definition of θ_u^c , the second equality follows by construction of p_u , and the second inequality follows because p_u is weakly increasing. For $x > u/(\beta\ell)$, we have

$$\phi(x, u) = g\left(\frac{\ell \cdot x}{u}\right) \leq g\left(\frac{\ell \cdot \bar{X}(\theta_u^c)}{u}\right) + g'\left(\frac{\ell \cdot \bar{X}(\theta_u^c)}{u}\right) \frac{\ell}{u} (x - \bar{X}(\theta_u^c)) = p_u(x),$$

where the inequality follows from the concavity of g . Thus, $\phi(\cdot, u) \leq p_u$ pointwise, establishing point (a) of DM.

By construction,

$$\text{supp}(H_u) = \{\underline{X}(\theta_u^c), \bar{X}(\theta_u^c)\} \subseteq [0, \theta_u^c] \cup \{\bar{X}(\theta_u^c)\} \subseteq \{x \in \Theta : p_u(x) = \phi(x, u)\},$$

establishing point (b) of DM.

Finally, $\mathbb{E}_{H_u}[p(x)] = p_u(\bar{X}(\theta_u^c))(1 - F(\theta_u^c))$ while

$$\begin{aligned} \mathbb{E}_F[p_u(x)] &= \int_{\theta_u^c}^1 \left[g\left(\frac{\ell \cdot \bar{X}(\theta_u^c)}{u}\right) + g'\left(\frac{\ell \cdot \bar{X}(\theta_u^c)}{u}\right) \frac{\ell}{u} (x - \bar{X}(\theta_u^c)) \right] dF(x) \\ &= p_u(\bar{X}(\theta_u^c))(1 - F(\theta_u^c)) + g'\left(\frac{\ell \cdot \bar{X}(\theta_u^c)}{u}\right) \frac{\ell}{u} \underbrace{\int_{\theta_u^c}^1 (x - \bar{X}(\theta_u^c)) dF(x)}_{=0} \\ &= p_u(\bar{X}(\theta_u^c))(1 - F(\theta_u^c)), \end{aligned}$$

where the last equality follows because, by definition, $\bar{X}(\theta_u^c) = \mathbb{E}_F[\tilde{\theta} | \tilde{\theta} \geq \theta_u^c]$, establishing point (c) of DM.

We can therefore conclude that $H_u \in \mathcal{H}(u)$. Additionally, by [Lemma 3](#), $\phi(\underline{X}(\theta_u^c), u) = 0 < \phi(\bar{X}(\theta_u^c), u)$. Since these properties hold for each $u \in (0, \beta\ell)$, they also hold for $u = \bar{u}$. Therefore, there exists a constrained-efficient binary market segmentation in which all buyers join only the submarket with the highest expected type, while the submarket with the lowest expected type remains inactive in equilibrium. This establishes Point (ii) of the proposition. \square

A.2. General characterization for constrained-efficient segmentations

This section provides a characterization for constrained-efficient market segmentations under [Assumption 1](#) and [Assumption 2](#) only.

From [Proposition 3](#), each market segmentation σ has an essentially-unique search equilibrium $(\tau_\sigma^*, u_\sigma^*)$ with $\tau_\sigma^*(\cdot) = \phi(\cdot, u_\sigma^*)$, where recall that

$$\phi(\nu, u) = g\left(\frac{\mathbb{E}_\nu[\lambda(\theta)\theta]}{u}\right) \mathbb{1}_{[\mathbb{E}_\nu[\lambda(\theta)\theta] > \frac{u}{\beta}]}.$$

Given a market segmentation σ and buyers' anticipated payoff u , define the ex-

ante total surplus as

$$\mathcal{S}(\sigma, u) := k \int_{\Delta(\Theta)} m(\phi(\nu, u)) \mathbb{E}_\nu[\theta] \sigma(d\nu),$$

and the ex-ante buyer's payoff as

$$\mathcal{U}(\sigma, u) = k \int_{\Delta(\Theta)} m(\phi(\nu, u)) \mathbb{E}_\nu[\lambda(\theta)\theta] \sigma(d\nu).$$

We can then express the planner's second-best problem (SB) as

$$\max_{\sigma} S(\sigma, u_{\sigma}^*).$$

The objective in (SB) is non-linear because the planner's choice of a market segmentation affects both what the buyers' believe about sellers' types within each submarket, and also how they trade off joining one submarket over another. Instead of directly solving (SB) directly, consider instead the following problem:

$$\max_{\substack{\sigma, \\ u \in [0, \beta \sup_{\theta \in \Theta} \lambda(\theta)\theta]}} \mathcal{S}(\sigma, u) \quad \text{s.t.} \quad \mathcal{U}(\sigma, u) \leq u. \quad (\text{BP})$$

The problem in (BP) decouples the two interdependent effects in (SB): holding fixed the anticipated payoff, the market segmentation now only affects beliefs, and holding fixed the market segmentation, the anticipated payoff affects the trade-offs across submarkets. The interdependence between these two effects is then captured by the constraint, $\mathcal{U}(\sigma, u) \leq u$, which ensures that buyers' anticipated payoffs are consistent with their expected gains from trade. As such, holding u fixed, the optimal choice of market segmentation in (BP) is akin to a (constrained) Bayesian Persuasion problem with a continuum of states [Kamenica and Gentzkow \(2011\)](#).

Proposition 6. *A market segmentation σ is constrained-efficient if and only if (σ, u_{σ}^*) is a solution to (BP).*

Proof of Proposition 6. Let σ' be a constrained-efficient market segmentation, i.e., σ' is a solution to (SB). Let (σ'', u'') be a solution to (BP).

Notice that for any market segmentation σ , the pair (σ, u_{σ}^*) satisfy the constraint in (BP), and thus, (BP) is a relaxation of (SB). Therefore, $\mathcal{S}(\sigma', u_{\sigma'}^*) \leq \mathcal{S}(\sigma'', u'')$.

Conversely, notice that $u'' > 0$; otherwise the constraint in (BP) would be violated. Additionally, we must have $\mathcal{U}(\sigma'', u'') = u''$; otherwise, the planner could lower the anticipated payoff by a small amount, which would improve the objective without violating the constraint. However, the constraint holding with equality implies $u'' = u_{\sigma''}^*$. Hence, $\mathcal{S}(\sigma', u_{\sigma'}^*) \geq \mathcal{S}(\sigma'', u_{\sigma''}^*) = \mathcal{S}(\sigma'', u'')$. \square

References

- Simon P Anderson and Régis Renault. Advertising content. *American Economic Review*, 96(1):93–113, 2006.
- Itai Arieli, Yakov Babichenko, Rann Smorodinsky, and Takuro Yamashita. Optimal persuasion via bi-pooling. *Theoretical Economics*, 18(1):15–36, 2023.
- Pak Hung Au and Mark Whitmeyer. Attraction versus persuasion: Information provision in search markets. *Journal of Political Economy*, 131(1):202–245, 2023.
- Dirk Bergemann, Benjamin Brooks, and Stephen Morris. The limits of price discrimination. *American Economic Review*, 105(3):921–957, 2015.
- David Blackwell. Equivalent comparisons of experiments. *The annals of mathematical statistics*, pages 265–272, 1953.
- David Blackwell and M.A. Girshick. *Theory of Games and Statistical Decisions*. Wiley, 1954.
- Simon Board and Jay Lu. Competitive information disclosure in search markets. *Journal of Political Economy*, 126(5):1965–2010, 2018.
- Gerard R Butters. Equilibrium distributions of sales and advertising prices. *The Review of Economic Studies*, 44(3):465–491, 1977.
- Xiaoming Cai, Pieter A Gautier, and Ronald P Wolthoff. Search frictions, competing mechanisms and optimal market segmentation. *Journal of Economic Theory*, 169:453–473, 2017.
- Michael Choi, Anovia Yifan Dai, and Kyungmin Kim. Consumer search and price competition. *Econometrica*, 86(4):1257–1281, 2018.
- Michael Choi, Kyungmin Kim, and Marilyn Pease. Optimal information design for search goods. In *AEA Papers and Proceedings*, volume 109, pages 550–556. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, 2019.
- Daniele Condorelli and Balazs Szentes. Buyer-optimal platform design. *Mimeo*, 2023.

- Peter A Diamond. Wage determination and efficiency in search equilibrium. *The Review of Economic Studies*, 49(2):217–227, 1982a.
- Peter A Diamond. Aggregate demand management in search equilibrium. *Journal of political Economy*, 90(5):881–894, 1982b.
- Mustafa Dogan and Ju Hu. Consumer search and optimal information. *The RAND Journal of Economics*, 53(2):386–403, 2022.
- Piotr Dworzak and Giorgio Martini. The simple economics of optimal persuasion. *Journal of Political Economy*, 127(5):1993–2048, 2019.
- Jan Eeckhout and Philipp Kircher. Sorting and decentralized price competition. *Econometrica*, 78(2):539–574, 2010a.
- Jan Eeckhout and Philipp Kircher. Sorting versus screening: Search frictions and competing mechanisms. *Journal of Economic Theory*, 145(4):1354–1385, 2010b.
- Matthew Elliott, Andrea Galeotti, Andrew Koh, and Wenhao Li. Market segmentation through information. *Available at SSRN 3432315*, 2021.
- Wei He and Jiangtao Li. Competitive information disclosure in random search markets. *Games and Economic Behavior*, 140:132–153, 2023.
- Arthur J Hosios. On the efficiency of matching and related models of search and unemployment. *The Review of Economic Studies*, 57(2):279–298, 1990.
- Ju Hu. Industry-optimal information in a search market. *Available at SSRN 3935299*, 2022.
- Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011.
- Kyungmin Kim. Endogenous market segmentation for lemons. *The RAND Journal of Economics*, 43(3):562–576, 2012.
- Kyungmin Kim and Philipp Kircher. Efficient competition through cheap talk: the case of competing auctions. *Econometrica*, 83(5):1849–1875, 2015.
- Andreas Kleiner, Benny Moldovanu, and Philipp Strack. Extreme points and majorization: Economic applications. *Econometrica*, 89(4):1557–1593, 2021.

- Anton Kolotilin, Timofiy Mylovanov, and Andriy Zapechelnyuk. Censorship as optimal persuasion. *Theoretical Economics*, 17(2):561–585, 2022.
- Stephan Laueremann. Asymmetric information in bilateral trade and in markets: An inversion result. *Journal of Economic Theory*, 147(5):1969–1997, 2012.
- Benjamin Lester, Ali Shourideh, Venky Venkateswaran, and Ariel Zetlin-Jones. Screening and adverse selection in frictional markets. *Journal of Political Economy*, 127(1):338–377, 2019.
- Wenhao Li. How to mismatch consumers. Technical report, Mimeo, 2023.
- David G Luenberger. *Optimization by vector space methods*. John Wiley & Sons, 1997.
- Chen Lyu. Information design for selling search goods and the effect of competition. *Journal of Economic Theory*, 213:105722, 2023.
- Teddy Mekonnen and Bobak Pakzad-Hurson. Competition, persuasion, and search. *arXiv preprint arXiv:2411.11183*, 2024.
- Teddy Mekonnen, Zeky Murra-Anton, and Bobak Pakzad-Hurson. Persuaded search. *arXiv preprint arXiv:2303.13409*, 2023.
- Guido Menzio. A theory of partially directed search. *Journal of political Economy*, 115(5):748–769, 2007.
- Espen R Moen. Competitive search equilibrium. *Journal of political Economy*, 105(2):385–411, 1997.
- James D Montgomery. Equilibrium wage dispersion and interindustry wage differentials. *The Quarterly Journal of Economics*, 106(1):163–179, 1991.
- Dale T Mortensen. The matching process as a noncooperative bargaining game. In *The economics of information and uncertainty*, pages 233–258. University of Chicago Press, 1982.
- Dale T Mortensen and Randall Wright. Competitive pricing and efficiency in search equilibrium. *International economic review*, 43(1):1–20, 2002.

- Michael Peters. Ex ante price offers in matching games non-steady states. *Econometrica: Journal of the Econometric Society*, pages 1425–1454, 1991.
- Barbara Petrongolo and Christopher A Pissarides. Looking into the black box: A survey of the matching function. *Journal of Economic literature*, 39(2):390–431, 2001.
- Christopher A Pissarides. Short-run equilibrium dynamics of unemployment, vacancies, and real wages. *The American Economic Review*, 75(4):676–690, 1985.
- Richard Rogerson, Robert Shimer, and Randall Wright. Search-theoretic models of the labor market: A survey. *Journal of economic literature*, 43(4):959–988, 2005.
- Gleb Romanyuk and Alex Smolin. Cream skinning and information design in matching markets. *American Economic Journal: Microeconomics*, 11(2):250–76, 2019.
- Hiroto Sato and Ryo Shirakawa. Persuasion in ordered search. *Available at SSRN 4483732*, 2023.
- Shouyong Shi. A directed search model of inequality with heterogeneous skills and skill-biased technology. *The Review of Economic Studies*, 69(2):467–491, 2002.
- Robert Shimer. The assignment of workers to jobs in an economy with coordination frictions. *Journal of Political Economy*, 113(5):996–1025, 2005.
- Nikhil Vellodi. Ratings design and barriers to entry. *Available at SSRN 3267061*, 2018.
- Randall Wright, Philipp Kircher, Benoît Julien, and Veronica Guerrieri. Directed search and competitive search equilibrium: A guided tour. *Journal of Economic Literature*, 59(1):90–148, 2021.