Competing for Consumer Inattention*

Geoffroy de Clippel† Kfir Eliaz‡ Kareen Rozen§

June 2013

Abstract

Consumers purchase multiple types of goods and services, but may be able to examine only a limited number of markets for the best price. We propose a simple model which captures these features, conveying some new insights. A firm's price can deflect or draw attention to its market, and consequently, limited attention introduces a new dimension of competition across markets. We fully characterize the resulting equilibrium, and show that the presence of partially attentive consumers improves consumer welfare as a whole. When consumers are less attentive, they are more likely to miss the best offer in each market; but the enhanced cross-market competition decreases average price paid, as leading firms try to stay under the consumers' radar.

*We are grateful to Phil Reny, Ran Spiegler and Bruno Strulovici for valuable suggestions and Xiaosheng Mu for excellent research assistance.

†Brown University, Department of Economics.

‡Tel Aviv University and the University of Michigan - Ann Arbor, Departments of Economics.

§Yale University, Department of Economics and Cowles Foundation. Rozen thanks the NSF for financial support under grant SES-0919955.

1 Introduction

Classic models of price competition assume that consumers have unlimited ability to track down the best deals. The wide array of goods and services in the marketplace casts doubt that this is a faithful description of the average consumer. With only limited attention to devote to finding cheaper substitutes, consumers may pay close attention to some purchases while neglecting to find the best price in others. This paper investigates the price and welfare implications of allocating limited attention across markets. Our simple model conveys some new insights: (i) a firm's price can deflect or draw attention to its market; and consequently, (ii) limited attention introduces a new dimension of competition across (even otherwise independent) markets.

We convey these insights in a simple framework, but they should remain important considerations in more general settings. Consumers in our model have unit demand for each of M different goods. To make point (ii) as starkly as possible, each consumer's utility is separable across goods, which ensures these markets would be independent if attention were unlimited. Reservation prices are assumed to be one for all consumers and all goods. Each good is offered by two sellers whose constant marginal cost is normalized to zero, and who set prices independently. For each market, consumers have a default seller who is interpreted as the most visible provider of that good or service. Consumers share the same default set of sellers, who are thought of as the *market leaders.* Confronted with market leaders' prices, consumers decide which markets to examine further, to see whether the competing firm (the *market challenger*, whose identity and price they do not know) offers a better deal. Consumers may have only limited attention to devote to comparison-shopping, with the ability to investigate at most $k \in \{0, \dots, M\}$ markets. The distribution of attention in the population is captured by a probability distribution $(\alpha_0, \dots, \alpha_M)$.

Our model captures the view that limited attention introduces an auditing component into consumption decisions. Given his budget of attention, a consumer uses what he knows (in this case, the price offered by market leaders) to decide which dimensions of his consumption decision are worthiest of further investigation. For instance, when buying groceries online, which items does a consumer buy from his saved list, and which does he check for better bargains? In a sense, a consumer's problem under limited attention is akin to that of maintenance scheduling in operations research: only a subset of items can be served, and those that are neglected may suffer from poor performance. For a consumer with limited attention, inspecting one market means overlooking another. The cost associated with this tradeoff is *endogenous*, equal to the expected equilibrium savings foregone by neglecting that other market.

Our setting is one of imperfect information, since consumers do not observe challengers' prices when allocating their attention. The analysis focuses on partially symmetric, perfect Bayesian Nash equilibria (henceforth *equilibria*). These preserve the symmetry of the model, with firms in the same position (as leaders or challengers) using the same pricing strategy. In that case, consumers expect the most savings to be found in markets with the most expensive leaders. Hence firms' profits may vary discontinuously with the leaders' prices, as consumers shift their attention between markets. A more standard form of discontinuity also arises when firms in a market quote the same price. Despite these discontinuities, we constructively establish that a partially symmetric equilibrium exists for any distribution of attention, and moreover, that only one such equilibrium exists. In this equilibrium, all firms employ atomless pricing strategies, but leaders systematically charge a wider range of prices than challengers. The support of the leaders' strategy has no gap. However, depending on the distribution of attention, challengers may avoid charging some intermediate prices. Constructing the unique equilibrium then requires an ironing procedure.

What is the equilibrium effect of (in)attention on consumer welfare? As might be expected, an increase in the proportion α_0 of fully inattentive consumers is detrimental. However, varying the distribution of partially attentive consumers has perhaps surprising implications. Any change in the distribution of attention which *decreases* the average level of attention (holding α_0 constant) is *beneficial*. This may seem unintuitive at first, since consumers inspecting fewer markets are more likely to miss the best deals. But this intuition does not take into account the countervailing effect of partial inattention on firms' behavior.[1] Consumers' limited capacity to search for better deals induces cross-market competition for their *in*attention: by lowering its price, a leader can increase the chance his market remains under the consumers' radar. The overall effect could, at least in theory, be determined by computing the consumer surplus directly using our expressions for the equilibrium strategies. Our argument follows a different route, taking advantage of the fact that total surplus remains constant and that firms' equilibrium profits turn out to be much simpler to calculate. We delve further into the mechanics of competition for inattention, exploring how the leaders' pricing strategy adjusts.

This paper proceeds as follows. In the next subsection we discuss how our paper fits within the literature. Section 2 presents the model. Section 3 presents the main results and their intuition, including how consumers allocate attention, the equilibrium characterization, and comparative statics with respect to partial attention. The constructive proof of the unique equilibrium is presented in Section 4. Concluding remarks, and possible directions for future research, are given in Section 5. Some proofs are relegated to the appendix.

Related literature

Our setting builds on the seminal literature on price dispersion (Salop and Stiglitz, 1977; Rosenthal, 1980; Varian, 1980), which explains observed variation in prices by introducing "captive" consumers who purchase from a randomly selected firm, without engaging in price comparisons. Among other differences with that literature, we consider multiple markets and introduce partially attentive consumers, which are driving forces behind our results. These and other features of our framework, such as the endogenous cost of neglecting a market and the asymmetric positions of firms, also depart from the standard approach taken in the search and rational inattention literatures.

[1] As an analogy, think of auctions under asymmetric information. Fixing the bids, first price gives a strictly higher profit than second price. However, this does not mean that equilibrium profits are necessarily higher with a first-price auction, as individuals' bidding behavior responds to the auction format.

In the search literature, consumers incur a fixed, exogenous cost of sampling prices of a product sold by multiple firms; classic references include Burdett and Judd (1983), where consumers decide in advance how many prices to simultaneously sample, or Stahl (1989), where consumers search sequentially. The rational inattention literature, pioneered by Sims (2003) and extended in Woodford (2009), introduces the use of entropy measures to model the exogenous cost of information processing. The consumers' dilemma in those literatures is whether to obtain any information, and if so, how much. In our approach, prices serve as a cue to determine which markets are worthiest of attention, which introduces an element of competition across sellers of different goods.

Market interaction between profit-maximizing firms and consumers with limited attention is, of course, more intricate than the stylized environment we analyze. Our model isolates an aspect of the feedback between consumer attention and firm behavior that has not been studied in the literature. One strand of this literature has focused on a different aspect of attention: when firms offer a multi-dimensional product, consumers may take only a subset of these dimensions into consideration. This approach is exemplified by Spiegler (2006), where a consumer samples one price dimension from each firm selling a product with a complicated pricing scheme (e.g., health insurance plans); Gabaix and Laibson (2006), where some consumers do not observe the price of an add-on before choosing a firm; Armstrong and Chen (2009), who extend the notion of "captive" consumers to those who always consider one dimension of a product but not another (say, price but not quality); and Bordalo, Gennaioli and Shleifer (2013), who study a duopoly model where firms decide on price and quality, taking into account that the relative weights consumers give to these attributes is determined endogenously by the choices of both firms. The above works study symmetric pricing equilibria for firms in a single market, with some differing implications for welfare. In Gabaix and Laibson (2006), for instance, prices increase as more consumers notice add-ons; while in Armstrong and Chen (2009), reducing the proportion of captive consumers reduces the incentive to offer low quality, but has an ambiguous effect on consumer welfare.

Taking a different approach to attention, Eliaz and Spiegler (2011a,b) formalize a model of competition over consumers who only consider a subset of available products. They abstract from prices and analyze firms who compete over market share only by offering a menu of products together with a payoff irrelevant marketing device (e.g., packaging). Consumers in their model are characterized by a preference relation and a consideration function, which determines, given firms' choices, whether a consumer pays attention only to its (exogenously determined) default firm or whether he also considers the competitor. They show that consumer welfare need not be monotonic in the amount of attention implied by the consideration function.

2 The model

We propose a simple model capturing the feature that consumers purchase multiple types of goods and services, but have the capacity to examine only a limited number of markets in search of the best price. The market for each good or service consists of two firms, a *leader* and a *challenger*, who compete in prices. All consumers know the market leaders' prices, but need to pay attention to a market to identify the challenger and learn his offer. Consumers differ in the number of markets to which they can pay attention. The leader in a market is interpreted as the most visible provider of the good or service, and is the default provider for a consumer who chooses not to allocate the time or capacity to search that market further.

There is a unit mass of consumers, each of whom desires at most one unit of any given good. For simplicity, we assume that the consumers' reservation price for each type of good is one. Letting M denote the number of markets (one per good), a consumer's utility from purchasing the bundle $(x_1, x_2, \dots, x_M) \in \{0,1\}^M$ at prices $(p_1, p_2, \dots, p_M)$ is $\sum_{m=1}^{M}(1-p_m)x_m$.

The distribution of attention in the consumer population is captured by a probability distribution $\alpha = (\alpha_0, \alpha_1, \dots, \alpha_M)$, where α_k is the proportion of consumers who can inspect up to k markets to find the best price. Consumers optimally decide which markets to inspect. If a consumer inspects a market, then he can choose whether to purchase from the market leader, the challenger, or not at all. If he does not inspect a market, then his only decision for that market is whether to purchase from its leading firm. The distribution of attention is common knowledge among firms. We assume throughout a positive measure of fully attentive consumers ($\alpha_M > 0$), inattentive consumers ($\alpha_0 > 0$), and partially attentive consumers ($\alpha_0 + \alpha_M < 1$).

The game unfolds over two periods. First, all firms independently set prices to maximize (expected) profit. We normalize marginal costs to zero, so realized profit is simply the product of the firm's price and its market share. Upon observing all the leaders' offers, consumers decide how to allocate their attention, and make their purchasing decisions, to maximize (expected) utility.

Equilibrium. Because consumers have only imperfect information when allocating their attention, the equilibrium notion applied is that of Perfect Bayesian equilibrium. We restrict attention throughout to *partially symmetric* equilibria where market leaders follow a common pricing strategy, as do market challengers. The leaders' strategy may differ from that of the challengers, and we do not impose restrictions on the consumers' strategies. We note that equilibrium existence is nontrivial, since firms' profits are discontinuous.[2]

Notation and definitions. The leaders' and challengers' pricing strategies are described by the cumulative distribution functions $F_\ell : \mathbb{R} \to [0,1]$ and $F_c : \mathbb{R} \to [0,1]$, respectively. A price p is said to be in the support of the pricing strategy F if $F(p+\varepsilon) - F(p-\varepsilon)$ is strictly positive for all $\varepsilon > 0$. A price p is said to be an atom of the strategy F if $\lim_{\varepsilon\to 0} F(p+\varepsilon) - F(p-\varepsilon) \neq 0$.

[2] Firms' payoffs exhibit two forms of discontinuity. The first, related to how a leader and a follower in a market share consumers when quoting the same price, appears in many models of competition. Existence in such cases follows from results by Dasgupta and Maskin (1986) or Reny (1999). The second form of discontinuity is related to how consumer attention is allocated across markets, and its impact on challengers' profits, when some leaders quote the same price. For each price he may quote, a challenger's profit is discontinuous over a continuum of leaders' prices, which prevents a direct application of Dasgupta and Maskin (1986). It also implies that challengers cannot secure themselves a positive payoff in the sense of Reny (1999). While alternative methods may be used to show existence, we provide a constructive proof that also establishes uniqueness.

We do not put *a priori* restrictions on the presence of atoms or gaps in the support of the pricing strategies.

3 Main results and intuitions

In this section, we first present our characterization of partially symmetric equilibria and some of the intuitions behind it, leaving the complete equilibrium analysis to Section 4. We then examine how the equilibrium and consumer welfare change with the distribution of attention among consumers.

3.1 Consumer attention and its implications

Suppose the leading firm in market i quotes a price p_i. A consumer's expected gain from inspecting that market is the expected savings from finding a cheaper price by the challenger, i.e.,

$$\int_0^{p_i} (p_i - q)dF_c(q). \tag{1}$$

Thus, the expected gain from inspecting market i is strictly positive if and only if the challenger quotes a price cheaper than p_i with positive probability. Let $S = \{i \mid F_c(p_i) > 0\}$ be the set of all such markets. Moreover, if $i \in S$ and market j's leader quotes a price strictly higher than p_i, then inspecting market j yields strictly higher expected savings than inspecting market i. The consumer's allocation of attention can thus be described as follows.

Proposition 1. *If a consumer with k units of attention inspects market i, then it is impossible to find k markets in S where the leader quotes a strictly higher price than p_i.*

To express firms' incentives, we must understand how a leader's price affects the probability with which consumers pay attention to its market. Through a series of results in Section 4, we show that leaders' prices are all distinct and the consumer has hope of finding a cheaper option in any market, for almost any prices quoted by leaders. In this case, Proposition 1 takes a simple form: the consumer inspects the k markets with the highest leader prices.

Market leaders. We can now compute the probability that a leader's market is paid attention to by a consumer with k units of attention, assuming that leader charges the price p and that all other market leaders follow the pricing strategy F_ℓ. Letting $x = F_\ell(p)$, we denote this probability by $\pi_k^\ell(x)$. Observe that his market receives attention from such a consumer if there are *no more than* $k-1$ *other markets* whose price turns out to be higher than p. Since the probability that another leader charges above p is $1-x$, we find that[3]

$$\pi_k^\ell(x) := \sum_{i=0}^{k-1} \binom{M-1}{i} x^{M-1-i}(1-x)^i. \tag{2}$$

As expected, $\pi_0^\ell(x) = 0$ and $\pi_M^\ell(x) = 1$. In addition, the probability of being inspected by a given consumer is increasing in his capacity for attention k, and increasing with one's price (as captured by x).

Market challengers. Consider a challenger's probability of selling to a consumer with k units of attention, assuming that he himself charges the price p and that all market leaders follow the pricing strategy F_ℓ. Letting $x = F_\ell(p)$, we denote this probability by $\pi_k^c(x)$. If a consumer is only partially attentive (that is, $k < M$), then $\pi_k^c(x)$ is not simply $1-x$, the *ex-ante* probability that the leader's price is higher than p. For the challenger, selling requires the consumer to pay attention to the market, an event whose probability is itself impacted by the leader's price. We may compute $\pi_k^c(x)$ as follows. The challenger has zero probability of making a sale if the leader in his market quotes a price strictly less than p. If the leader quotes a price $q > p$, then the consumer will purchase from the challenger so long as he inspects the market, which occurs with probability $\pi_k^\ell(F_\ell(q))$. Integrating over the possible prices of the market leader, the desired probability is given by $\int_p^\infty \pi_k^\ell(F_\ell(q))dF_\ell(q)$.

[3] This amounts to having at most $k-1$ "successes" in $M-1$ trials that are i.i.d., where the probability of "success" (which means finding a price higher than p) is $1-F_\ell(p)$.

This probability depends only on $x = F_\ell(p)$ and not the entire distribution F_ℓ, as can be seen using the change of variables $t = F_\ell(q)$:

$$\pi_k^c(x) := \int_x^1 \pi_k^\ell(t)dt. \tag{3}$$

As expected, $\pi_0^c(x) = 0$ and $\pi_M^c(x) = 1 - x$. In addition, the probability of selling to a given consumer is increasing in his capacity for attention k, and decreasing with the probability x that the leader's price is better.

3.2 Equilibrium characterization

It will be helpful to define the total probability that a leader's market draws attention if he charges a price p, and the total probability that a market challenger sells if he charges a price p. Recalling that α is the distribution of attention among consumers, and letting $x = F_\ell(p)$, those probabilities are

$$\Pi_\ell(x) := \sum_{k=1}^{M} \alpha_k \pi_k^\ell(x) \text{ and } \Pi_c(x) := \sum_{k=1}^{M} \alpha_k \pi_k^c(x),$$

respectively. Since there is a positive measure of partially attentive consumers, Π_ℓ is strictly increasing and Π_c is strictly decreasing; hence their inverses Π_ℓ^{-1} and Π_c^{-1} are well-defined.

Deriving indifference conditions. In Section 4, we show the following. First, if an equilibrium exists, then α_0 is the lowest price in the support of both the leaders' and challengers' strategies. Second, both leaders' and challengers' pricing strategies must be atomless. Third, the leaders' strategy has full support over the interval $[\alpha_0, 1]$, while the challenger's highest price $\overline{p}_c$ must be strictly smaller than 1.

These equilibrium properties imply that a leader's profit from charging each price in $[\alpha_0, 1]$ must equal his profit from charging the price 1. This profit is simply α_0, given that only fully inattentive consumers would purchase from him. In other words, for any $p \in [\alpha_0, 1]$,

$$p\Big(1 - \Pi_\ell\big(F_\ell(p)\big) + \Pi_\ell\big(F_\ell(p)\big)\big(1 - F_c(p)\big)\Big) = \alpha_0, \tag{4}$$

since the leader sells at the price p either when a consumer does not pay attention, or when he pays attention but the challenger's price is higher.

Similarly, a challenger's profit from each price in its support must equal its profit from quoting α_0. This profit is given by $\alpha_0 \Pi_c(0)$, which in turn equals $\alpha_0 EA(\alpha)/M$, where

$$EA(\alpha) := \sum_{k=1}^{M} \alpha_k k$$

is the expected level of attention in the consumer population. Indeed, because the leaders' strategy is atomless and prescribes only prices above α_0, the challenger is sure to sell to consumers who pay attention; and given that market leaders all use the pricing strategy F_ℓ, there is a k out of M chance that his market leader's price will be among the k-highest.[4] Therefore, for any price p in the support of F_c, it must be that

$$p\Pi_c\big(F_\ell(p)\big) = \frac{\alpha_0 EA(\alpha)}{M}. \tag{5}$$

For any price in the support of a challenger's strategy, the leader's strategy is defined by the indifference condition (5); for all other prices, it is defined by the indifference condition (4), as a function of the (constant) level of F_c. In other words,

$$F_\ell(p) = \begin{cases} \Pi_c^{-1}\left(\frac{\alpha_0 EA(\alpha)}{Mp}\right) & \text{for all } p \text{ in the support of } F_c, \\ \Pi_\ell^{-1}\left(\frac{p-\alpha_0}{pF_c(p)}\right) & \text{for all other } p \in [\alpha_0, 1]. \end{cases} \tag{6}$$

The challenger's strategy is also defined by the indifference condition (4) for any price in its support. Solving for F_c in (4) and applying the definition of

[4] This can also be seen by applying the Euler integral $\int_0^1 t^{a-1}(1-t)^{b-1} = \frac{(a-1)!(b-1)!}{(a+b-1)!}$ in the definition of $\pi_k^c(0)$ to show that it simplifies to k/M.

F_ℓ above, we see that for each price in the support of the challengers' pricing strategy, F_c must coincide with the function $\tilde{F}_c$ defined by

$$\tilde{F}_c(p) := \frac{p - \alpha_0}{p\Pi_\ell\Big(\Pi_c^{-1}\big(\frac{\alpha_0 EA(\alpha)}{Mp}\big)\Big)}, \text{ for all } p \in [\alpha_0, 1]. \tag{7}$$

Which prices does a challenger charge? The difficulty lies in knowing the support of the challengers' strategy, since Figure 1 demonstrates that $\tilde{F}_c$ may be *nonmonotonic* in p without further restrictions on the distribution of attention among consumers. That is, if we try to construct a putative equilibrium where the challengers' support is the entire interval $[\alpha_0, p_c]$, for some $p_c < 1$, then F_c would be given by $\tilde{F}_c$, which may not be a valid distribution function. If an equilibrium exists, then any nonmonotonicity in $\tilde{F}_c$ must be "ironed" by introducing one or more gaps in the support of the challengers' strategy.

Due to the absence of atoms, F_c must be continuous. Hence any single gap in F_c must be an interval between two prices whose $\tilde{F}_c$ values coincide. In Figure 1, for instance, a gap cannot start at a price lower than p_1. On the other hand, there is a range of prices larger than p_1 which can serve as the leftmost endpoint of a gap. Remember that the leaders' pricing strategy F_ℓ is defined piecewise in (6) according to the challengers' support. Can $\tilde{F}_c$ be ironed in a way that ensures F_ℓ is increasing and atomless, as we know it must be? These requirements turn out to be unrestrictive: F_ℓ satisfies them whenever $\tilde{F}_c$ is ironed in the continuous manner described above.

Thus there are infinitely many ways to construct valid distribution functions F_ℓ and F_c which leave the leaders and challengers indifferent over all prices in their respective supports. However, there is a *unique* way to iron $\tilde{F}_c$ that yields *equilibrium* distribution functions F_ℓ and F_c. Using any other approach, the challenger has a profitable deviation to prices outside of his support.

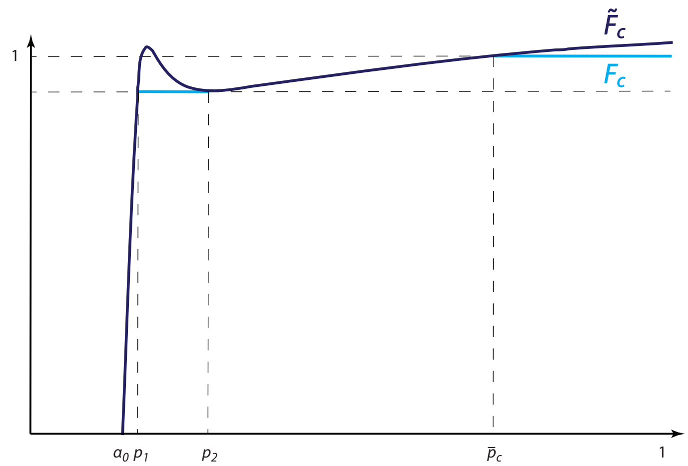

FIGURE 1: The construction of F_c in an example where $\tilde{F}_c$ is not increasing.

Theorem 1. *For any distribution of attention α, there exists a unique partially symmetric equilibrium. The challengers' pricing strategy F_c is atomless and given by*

$$F_c(p) = \min_{\tilde{p} \in [p,1]} \tilde{F}_c(\tilde{p}), \text{ for all } p \in [\alpha_0, \overline{p}_c], \tag{8}$$

where $\overline{p}_c \in (\alpha_0, 1)$ is the smallest price for which the above expression equals one, and $\tilde{F}_c$ is given by Equation (7). The leaders' pricing strategy F_ℓ has full support on $[\alpha_0, 1]$, is atomless, and given by Equation (6).

Theorem 1 is proved in Section 4. There we provide a complete equilibrium analysis, covering some important steps (e.g., ruling out the presence of atoms, characterizing the support of the leader) that have been glossed over in this section when deriving necessary equilibrium conditions. Moreover, we resolve the question of existence by showing that the construction is well-defined and indeed yields an equilibrium.

To state the characterization of F_c a bit differently, note that among all pricing strategies which lie below the graph of $\tilde{F}_c$, the challengers' strategy is the one which is pointwise highest. Hence it prescribes the "cheapest" price distribution among those, in the sense of first-order stochastic dominance. Graphically, this means $\tilde{F}_c$ must be ironed as illustrated in Figure 1, by starting any gap at the smallest possible price while still preserving continuity. For some intuition, suppose the challengers' pricing strategy F_c excludes a price p from its support, and yet $F_c(p) > \tilde{F}_c(p)$. Since p is outside the challengers' support, the F_ℓ is constructed so as to generate the leaders' equilibrium profit when charging p. As such, the more likely are challengers to charge *below* p, the more likely are leaders to charge *above* p. This may at first seem counterintuitive. However, the probability with which other leaders charge above p must increase, precisely to draw less attention to the market of a leader who charges p, and is thus more likely to be underbid by the challenger. The problem is that this twists the incentives of a challenger, who would then prefer to charge p rather than a smaller price in his support.

The presence of a gap in the challengers' strategy depends on the way attention is distributed among consumers. For any attention distribution α, the distribution of *partial attention* is $a(\alpha) = (\frac{\alpha_1}{1-\alpha_0}, \dots, \frac{\alpha_M}{1-\alpha_0})$. This is simply α conditioned on consumers being at least partially attentive, that is, $k \geq 1$. The lack of monotonicity in Figure 1 can be attributed to having multiple peaks in the partial attention distribution. Gaps can be ruled out when, given the proportion of consumers with attention span k and the proportion with attention span $k+2$, there are sufficiently many consumers falling in between. More formally, the partial attention distribution is *log-concave* if $\alpha_k^2 \geq \alpha_{k-1}\alpha_{k+1}$ for each $k \in \{2, \dots, M-1\}$, or equivalently, the likelihood ratio α_{k+1}/α_k is decreasing in k. Note that this is trivially satisfied when there are only two markets, and is implied whenever the entire attention distribution is log-concave. When partial attention has this feature, the form of the equilibrium pricing strategies simplifies.

Theorem 2. *When $\tilde{F}_c$ is strictly increasing, the challengers' pricing strategy F_c has full support on $[\alpha_0, \bar{p}_c]$ and the leaders' pricing strategy F_ℓ simplifies to*

$$F_\ell(p) = \max\left\{\Pi_c^{-1}\Big(\frac{\alpha_0 EA(\alpha)}{Mp}\Big)\ ,\ \Pi_\ell^{-1}\Big(1 - \frac{\alpha_0}{p}\Big)\right\}$$

for all $p \in [\alpha_0, 1]$. A sufficient condition for $\tilde{F}_c$ to be strictly increasing is log-concavity of the partial attention distribution.

Theorem 2 is proved in the appendix. Many distributions (and their truncations) satisfy log-concavity. For example, the property is satisfied by a positive binomial distribution, where consumers start with M units of attention but can lose up to $M-1$ of them due to independent, exogenously occurring emergencies (e.g., the consumer's washing machine breaks down, his child gets the flu, his boss asks for overtime, etc.). We note that gaps can also be ruled out under other assumptions on partial attention, such as when the distribution $a(\alpha)$ is increasing (that is, $a_k(\alpha) \leq a_{k+1}(\alpha)$ for each k).[5]

3.3 The comparative statics of attention

A natural question that arises from our analysis is how partial attention affects consumer surplus. One might think that paying more attention to markets would keep prices down. However, partially attentive consumers introduce a form of competition across otherwise independent markets, and improve consumer welfare as a whole.

Theorem 3. *Consider two distributions of attention α and $\hat{\alpha}$ which share the same proportion of fully inattentive consumers ($\alpha_0 = \hat{\alpha}_0$). Then consumer welfare is higher under α than $\hat{\alpha}$ if, and only if, the expected level of attention under α is lower than under $\hat{\alpha}$.*

Proof. By Corollary 1 in Section 4, a leader's equilibrium expected profit is equal to the proportion of fully inattentive consumers, and is thus the same under both α and $\hat{\alpha}$. By Corollary 2 in Section 4, a challenger's equilibrium expected profit is equal to the proportion of fully inattentive consumers, multiplied by the expected level of attention, divided by M. Hence producer surplus is lower under α than $\hat{\alpha}$ if, and only if, the expected level of attention under α is lower than under $\hat{\alpha}$. The result then follows from the fact that total surplus remains constant (equal to M). ∎

[5]We show in the appendix that gaps can be ruled out when $\Pi_c(0) - \Pi_c(x)$ is strictly log-concave. While $\Pi_c(0) - \Pi_c(x)$ can be written as the sum of log-concave functions, log-concavity is not necessarily preserved by aggregation. We show that log-concavity is preserved if the sequence $\beta_1 = \alpha_M$, $\beta_k = \beta_{k-1} + \sum_{i=1}^{k} \alpha_{M-i+1}$ is log-concave. This is implied not only when $a(\alpha)$ is log-concave, but also, for instance, when $a(\alpha)$ is increasing.

Neither fully attentive nor fully inattentive consumers generate competition for inattention. While fully attentive consumers do generate within-market competition, fully inattentive consumers are simply captive to market leaders. As might be expected, increasing the proportion α_0 of captive consumers has a negative effect on consumer surplus.[6] At the opposite end of the attention spectrum, Theorem 3 means that making fully attentive consumers less attentive (e.g., shifting mass from α_M to α_k, for some $k > 0$) benefits consumers as a whole.

To gain some intuition for Theorem 3, remember that in equilibrium, leaders are willing to quote prices that are more expensive than what a challenger would ever charge. When charging such a price p, a leader's profit, given by $p(1-\Pi_\ell(F_\ell(p)))$, relies on *not* drawing too much consumer attention. Suppose partial attention decreases. If the other leaders' pricing strategy were to remain unchanged, then the leader's profit from quoting p would rise above α_0. Yet competition implies that no leader can make a profit that large. Hence the likelihood of having other leaders quote prices smaller than p must go up, so that the leader quoting p "sticks out" with sufficient probability.

The pricing effects of a change in partial attention may be more ambiguous for lower prices, as leaders become competitive against the challengers. Building on the insight from Theorem 2, we focus on cases where $\tilde{F}_c$ is strictly increasing and show that the leaders' pricing strategies are comparable under first-order stochastic dominance when the change in partial attention can be ranked in the monotone likelihood ratio order. The distribution $a(\hat{\alpha})$ of partial attention dominates another distribution $a(\alpha)$ in the *monotone likelihood ratio order (MLR)* if $a_k(\hat{\alpha})/a_k(\alpha)$ is increasing in $k \in \{1, \ldots, M\}$. The MLR ordering has a long tradition in economics, starting with Milgrom (1981), and is known to be stronger than first-order stochastic dominance.

Theorem 4. *Let α and $\hat{\alpha}$ be two attention distributions with $\alpha_0 = \hat{\alpha}_0$, and for which the partial attention distributions are log-concave. If $a(\hat{\alpha})$ dominates $a(\alpha)$ in the MLR order, then market leaders' equilibrium prices are higher*

[6] Increasing α_0 at the expense of reducing $(\alpha_1, \ldots, \alpha_M)$ by the infinitesimal amounts $(\varepsilon_1, \ldots, \varepsilon_M)$ has a total effect on producer surplus of $\sum_{i=1}^{M} \varepsilon_i(M + EA(\alpha) - \alpha_0 i) > 0$.

under $\hat{\alpha}$ than under α, in the sense of first-order stochastic dominance.

More generally, Theorem 4 remains true when replacing the log-concavity requirement with any conditions on α and $\hat{\alpha}$ guaranteeing that the challengers' strategy has no gap (e.g., as in footnote 5). To see why the result holds, remember that Π_ℓ and Π_c (the probabilities that a leader's market receives attention and that a challenger makes a sale) depend on the attention distribution. In what follows, Π_ℓ and Π_c correspond to the attention distribution α, while $\hat{\Pi}_\ell$ and $\hat{\Pi}_c$ correspond to the attention distribution $\hat{\alpha}$. Recall from Theorem 2 that the probability $F_\ell(p)$ that a leader charges a price lower than p under attention distribution α is simply

$$\max\Big\{\Pi_c^{-1}\big(\frac{\alpha_0 EA(\alpha)}{Mp}\big)\ ,\ \Pi_\ell^{-1}\big(1-\frac{\alpha_0}{p}\big)\Big\}, \tag{9}$$

when the partial attention distribution is log-concave. An analogous expression describes the probability $\hat{F}_\ell(p)$ that a leader charges a price lower than p under attention distribution $\hat{\alpha}$. This is illustrated in Figure 2.

We now show, as illustrated, that each of the two expressions on the right-hand side of (9) shifts *downwards* when consumer attention increases from α to $\hat{\alpha}$. Consequently, market leaders charge first-order stochastically higher prices when attention increases.

The downward shift is easiest to see for the second expression in (9), which corresponds to the intuition given above, and in fact holds for any first-order stochastic increase in partial attention. Remember that $\Pi_\ell(x)$ is a weighted average, under α, of the probability $\pi_k^\ell(x)$ that a leader's market is inspected by a consumer with k units of attention, when there is probability x that other market leaders are cheaper. The higher a consumer's attention level k, the higher is this probability $\pi_k^\ell(x)$. So a first-order increase in the attention distribution means that given any x, a leader now faces a higher total probability of drawing consumers' attention. The downward shift in the second expression in (9) immediately follows.

We now consider the first expression in (9). Let x and $\hat{x}$ satisfy $\Pi_c(x) = \alpha_0 EA(\alpha)/Mp$ and $\hat{\Pi}_c(\hat{x}) = \hat{\alpha}_0 EA(\hat{\alpha})/Mp$. Showing $x > \hat{x}$ amounts to proving

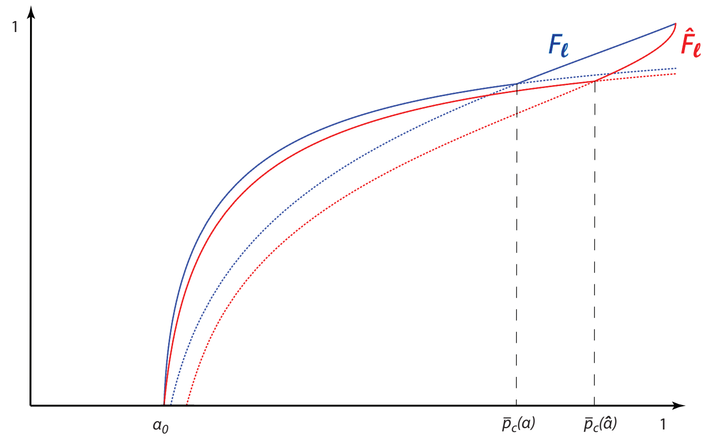

FIGURE 2: Comparative Statics on F_ℓ. The solid curves depict the market leaders' pricing strategies, which are the upper envelope of the corresponding dotted curves. The top (blue) pair of curves corresponds to attention distribution α, while the bottom (red) pair corresponds to $\hat{\alpha}$.

$\hat{\Pi}_c(x) < \hat{\Pi}_c(\hat{x})$, as the probability a challenger makes a sale, $\hat{\Pi}_c$, is decreasing in the probability x that his leader is cheaper. Consider the ratio of these expressions, which we can multiply and divide by $\Pi_c(x)$, and simplify using the definitions of x and $\hat{x}$:

$$\frac{\hat{\Pi}_c(x)}{\hat{\Pi}_c(\hat{x})} = \frac{\hat{\Pi}_c(x)}{\Pi_c(x)} \frac{\Pi_c(x)}{\hat{\Pi}_c(\hat{x})} = \frac{\hat{\Pi}_c(x)}{\Pi_c(x)} \frac{EA(\alpha)}{EA(\hat{\alpha})}.$$

This ratio is smaller than 1 if and only if $\hat{\Pi}_c(x)/\Pi_c(x) < EA(\hat{\alpha})/EA(\alpha)$. Observe that $EA(\hat{\alpha})/EA(\alpha)$ is the value of the ratio $\hat{\Pi}_c/\Pi_c$ evaluated at zero. Hence $x > \hat{x}$ if and only if this ratio is below its value at zero. We show in the appendix that the assumptions on α and $\hat{\alpha}$ guarantee this property. In particular, if the partial attention distribution $a(\hat{\alpha})$ dominates $a(\alpha)$ in the MLR order, then the ratio $\hat{\Pi}_c/\Pi_c$ is decreasing in x, and strictly so at zero.[7]

Changes in partial attention have a more ambiguous effect on the challengers' pricing strategy. Since consumer welfare increases when there is less

[7]This sufficient condition is not necessary, as the attention distributions used for Figure 2 do not have the MLR property but have the feature that $\hat{\Pi}_c/\Pi_c$ decreases.

attention, it is clear that challengers cannot increase their prices by too much. When there are just two markets, log-concavity of the partial attention distribution is trivially satisfied, and MLR-dominance reduces to first-order stochastic dominance. In that case, one can show that *both* leaders' and challengers' prices decrease when partial attention decreases. More generally, however, it is unclear whether the challengers' strategy shifts according to first-order stochastic dominance.

4 Complete equilibrium analysis

Building on the characterization of consumer attention in Proposition 1, we first develop a series of necessary conditions on firms' equilibrium pricing strategies that uniquely pin down the equilibrium, if one exists. We then resolve the matter of existence by checking that the construction works.

4.1 Necessary conditions

We begin with a useful observation about the supports of the challengers' and leaders' strategies.

Proposition 2. *The lowest price in the support of the pricing strategies of market leaders and challengers coincide, and is greater than or equal to* α_0*. The highest prices in the supports of the pricing strategies of market leaders and challengers are smaller than or equal to one.*

Proof. A market leader is sure to sell to inattentive consumers, even when charging the reservation price of 1. He can thus guarantee himself a profit of at least α_0. Any price below α_0 or above 1 generates a profit strictly less than α_0. Hence any strategy for which $F_\ell(1) < 1$ or $F_\ell(p) > 0$, for some $p < \alpha_0$, is strictly dominated.

Suppose that $F_c(p) > 0$ for some $p < \alpha_0$. Since $F_\ell(\alpha_0)$ must be zero, the challenger sells to the same set of consumers when charging a price $p' \in (p, \alpha_0)$ instead of p. Hence a pricing strategy for which $F_c(p) > 0$ for some $p < \alpha_0$

cannot be a best response against F_ℓ. Suppose now that $F_c(1) < 1$. Any price above 1 does not yield a sale, as it is higher than the consumers' reservation price. Any price $p < \alpha_0$ constitutes a profitable deviation for the challenger, because at least the fully attentive consumers will purchase from him.

Finally, suppose that the lowest price p in the support of one firm's strategy is strictly smaller than the lowest price p' in the support of his competitor in the same market. Let F be the pricing strategy corresponding to p. Suppose the firm using F deviates to a pricing strategy that has an atom equal to $F(p'-\varepsilon)$ at price $p' - \varepsilon$, and coincides with F for all higher prices. This deviation is strictly profitable for a challenger since there is positive probability consumers pay attention to his market. It is also strictly profitable for a leader, who underbids the challenger even if the deviation draws more attention. ∎

We next argue that F_ℓ is atomless. If leaders have an atom at a price strictly above the lowest price $\underline{p}$ in their support, then each *leader* could profitably deviate by moving mass from this price to one which is "slightly" below it. This small price decrease is more than compensated by the decreased attention to the leader's market. However, if the leaders' atom is on $\underline{p}$, we must distinguish between two cases. If the challenger's strategy does not have an atom at $\underline{p}$, or if some consumers favor the leader in case of a tie at $\underline{p}$, then the *challenger* could profitably deviate by shifting weight to prices slightly below $\underline{p}$. Otherwise, a *leader* can profitably deviate for the same reasons as given above.

Proposition 3. *The leaders' pricing strategy F_ℓ is atomless.*

Proof. Let $\underline{p}$ be the smallest price in the support of F_ℓ, and suppose F_ℓ has an atom at $p \in (\underline{p}, 1]$. For any small $\varepsilon > 0$, consider the alternate pricing strategy for the leader which equals $F_\ell(p)$ for all $q \in [p - \varepsilon, p]$, and coincides with F_ℓ elsewhere. This deviation has two opposite effects on the leader's profit. There is a negative effect from selling at a price $p - \varepsilon$ compared to those prices $q \in (p - \varepsilon, p]$. This loss is of order ε and can be made as small as desired by decreasing ε. In view of Proposition 1, there is also a positive effect from the decrease in attention when charging $p - \varepsilon$ rather than a price in $(p - \varepsilon, p]$. This gain admits a strictly positive lower bound that is independent of ε. To see this, consider the event that all other leaders charge p and that one's challenger has a price less than p, which occurs with probability $F_c(p)(F_\ell(p)-F_\ell(p^-))^{M-1} > 0$ since p is an atom and $F_c(p) > 0$ by Proposition 2. For the fraction $1-\alpha_0-\alpha_M$ of partially attentive consumers, if the deviator charges $p-\varepsilon$ they surely do not pay attention to his market, while if he charges p the probability of drawing attention is strictly positive (and independent of ε). Hence this deviation is strictly profitable for $\varepsilon > 0$ small enough.

Suppose now that F_ℓ has an atom at $\underline{p}$, which is also the lowest price in the support of F_c by Proposition 2. For any small $\varepsilon > 0$, consider the alternate pricing strategy for the challenger which equals $F_c(\underline{p}+\varepsilon)$ for all $q \in [\underline{p}-\varepsilon, \underline{p}+\varepsilon]$, and coincides with F_c elsewhere. This deviation has two opposite effects on the challenger's profit. There is a negative effect from selling at a lower price $\underline{p}-\varepsilon$ compared to those $q \in (\underline{p}, \underline{p}+\varepsilon]$; this results in a decrease in profit of no more than $2\varepsilon F_c(\underline{p}+\varepsilon)$. There is also a positive effect occurring in the event that the market draws attention when the leader's price is $\underline{p}$, which occurs with strictly positive probability (independent of ε). In this event, the deviation yields a sale at the price $\underline{p}-\varepsilon$ with probability $F_c(\underline{p}+\varepsilon)$, while the original strategy yields a sale at the price $\underline{p}$ with probability $F_c(\underline{p})\beta$, where $\beta \in [0,1]$ is the probability consumers purchase from the challenger when there is a tie at $\underline{p}$. The challenger's alternate strategy is a profitable deviation for $\varepsilon > 0$ small enough if either $F_c(\underline{p}) = 0$ (that is, the challenger does not have an atom at $\underline{p}$), or there is an atom at $\underline{p}$ and $\beta < 1$. To conclude the proof, suppose that both F_c and F_ℓ have an atom at $\underline{p}$ and $\beta = 1$. In that case, the same reasoning as in the first paragraph (applied with $p = \underline{p}$) shows that the market leader has a profitable deviation from F_ℓ, since $\beta = 1$ implies the challenger wins in case of a tie at $\underline{p}$. ∎

Our next result is concerned with the highest prices firms could charge. Challengers, who make their profit by underbidding their market leader, certainly would not charge more than a leader's highest price. We show, furthermore, that challengers charge strictly less. Since we have not yet ruled out the possibility that F_c has an atom at its highest price (or elsewhere), the strict ranking of highest prices is helpful to derive the leaders' highest price. Whenever a leader charges his highest price, any consumer who is at least partially attentive will inspect his market, and find a cheaper alternative, with probability one. As such, the leader may as well take full advantage of the remaining consumers' inattention, by charging all the way up to their reservation price.

Proposition 4. *The highest price in the support of F_c is strictly smaller than the highest price in the support of F_ℓ, which is one.*

Proof. Let $\overline{p}_\ell$ ($\overline{p}_c$) be the highest price in the support of F_ℓ (respectively, F_c). Since F_ℓ is atomless, there exists $\varepsilon > 0$ small enough that the probability a leader charges more than $\overline{p}_\ell - \varepsilon$ is strictly smaller than α_0. Thus the challenger's profit from charging any price above $\overline{p}_\ell - \varepsilon$ is strictly smaller than the profit obtained by charging α_0, given that he cannot affect the attention to his market. Since the challenger would have a profitable deviation if $F_c(\overline{p}_\ell - \varepsilon) < 1$, we conclude that $\overline{p}_c < \overline{p}_\ell$, .

We now show that the leaders' highest price is one. If $\overline{p}_\ell < 1$, then for each $\varepsilon > 0$, consider the alternate pricing strategy for a market leader which equals $F_\ell(\overline{p}_\ell - \varepsilon)$ for all $q \in [\overline{p}_\ell - \varepsilon, 1)$ and coincides with F_ℓ elsewhere. For ε small enough, $\overline{p}_\ell - \varepsilon$ is larger than the highest price in the support of F_c. The leader's gain from charging 1 instead of p is at least $\alpha_0(1 - \overline{p}_\ell)$, for each $p \in [\overline{p}_\ell - \varepsilon, \overline{p}_\ell)$, since fully inattentive consumers buy from the leader as long as his price is below their reservation level. The leader's loss from this deviation is proportional to the increase in probability of having partially attentive consumers check his market (thereby finding a cheaper price). For any $p \in [\overline{p}_\ell - \varepsilon, \overline{p}_\ell)$, this increase in probability converges to zero as ε decreases: F_ℓ atomless and ε small implies that it is almost certain that when charging $p \in [\overline{p}_\ell - \varepsilon, \overline{p}_\ell)$ his market was already being checked by these consumers. The expected change in profit is obtained by integrating gains minus losses over $p \in [\overline{p}_\ell - \varepsilon, \overline{p}_\ell)$. For ε small enough, the integrand is positive and hence the deviation is profitable. ∎

The above results now allow us to derive the leaders' equilibrium profit.

Corollary 1. *The leaders' equilibrium profit is α_0.*

Proof. Since the price 1 belongs to the support of F_ℓ, each market leader's equilibrium profit is equal to the profit made when charging 1. Because F_ℓ is atomless by Proposition 3, and because the highest price in the support of challengers is strictly less than 1 by Proposition 4, this profit must equal α_0. Only fully inattentive consumers buy from the market leader at that price, since there is probability one that all other consumers will check his market and purchase from the challenger. ∎

It now becomes possible to identify the common lowest price of challengers and leaders.

Proposition 5. *The lowest price in the support of both F_ℓ and F_c is α_0.*

Proof. We know that F_ℓ and F_c share a common lowest price $\underline{p} \geq \alpha_0$. Suppose by contradiction that $\underline{p} > \alpha_0$. Consider a deviation where the leader charges $(p+\alpha_0)/2$ with probability one. In this case, the leader sells to all consumers, whether or not they pay attention to his market. This delivers a profit of $(p+\alpha_0)/2$. Since equilibrium profit is α_0 by Corollary 1, the deviation is strictly profitable. ∎

We can now also derive the equilibrium profit of the challengers.

Corollary 2. *The challengers' equilibrium profit is $\alpha_0 EA(\alpha)/M$.*

Proof. By Proposition 5, α_0 is the lowest price in the support of both leaders and challengers. A challenger's equilibrium profit must thus equal its profit from quoting α_0. Since the leaders' strategy is atomless by Proposition 3, quoting α_0 gives profit $\alpha_0\Pi_c(0)$. By symmetry of the leaders' strategies, we know $\Pi_c(0) = EA(\alpha)/M$. ∎

The following result rules out atoms for the challenger. Of course, Propositions 3 and 6 imply that no firm can use a pure strategy in equilibrium.

Proposition 6. *The challengers' pricing strategy F_c is atomless.*

Proof. Suppose that F_c has an atom at some price $p > \alpha_0$. We begin by pointing out that there cannot exist $\varepsilon > 0$ for which $F_\ell(p+\varepsilon) - F_\ell(p) = 0$. Otherwise, F_ℓ has a gap in its support to the right of p, and the challenger could profitably deviate by shifting his atom from p to $p + \varepsilon$.

Consider an alternate strategy for the leader which equals $F_\ell(p+\varepsilon)$ for $q \in [p-\varepsilon, p+\varepsilon]$ and is given by F_ℓ elsewhere. For each $q \in [p-\varepsilon, p+\varepsilon]$, the only loss associated with this deviation is the decrease in price, which is at most 2ε. Among the various gains in profit from switching is the increased probability of selling by underbidding the challenger when the market is examined. Fix $p^* \in (\alpha_0, p)$. In view of Proposition 5, for any $\varepsilon < p - p^*$ the probability of the market drawing attention when the leader charges $p - \varepsilon$ is bounded below by a positive number that depends only on p^* and F_ℓ (not ε). Conditional on the market drawing attention, the increase in the probability of selling is bounded below by a number that is positive and independent of ε, due to the atom. Hence, for ε small enough, this deviation is strictly profitable for the leader, a contradiction. We conclude F_c is atomless for prices above α_0.

Finally, suppose by contradiction that F_c has an atom at α_0. Since α_0 also belongs to the support of F_ℓ, the leader must get a profit α_0 by charging any price $p \in (\alpha_0, \alpha_0 + \varepsilon)$. However, for any such price there is probability larger than $\alpha_M F_c(\alpha_0) > 0$ that the leader does not sell. Hence the profit from any such p is bounded away from α_0 for small ε, a contradiction. ∎

We now examine whether firms necessarily use *strictly* increasing strategies. While for market leaders the answer is a clear yes, for market challengers the answer depends on the distribution of consumer types. This contrasts with the previous literature on competition with mixed strategies over prices, in which all firms use strictly increasing cumulative distribution functions.

Proposition 7. *The leaders' strategy F_ℓ cannot have any gaps in its support.*

Proof. Suppose that F_ℓ has a gap in its support, that is, F_ℓ is constant over an interval inside $[\alpha_0, 1]$. Consider then p' and p'' with $F_\ell(p') = F_\ell(p'')$ such that for all $\varepsilon > 0$, $F_\ell(p') > F_\ell(p' - \varepsilon)$ and $F_\ell(p'' + \varepsilon) > F_\ell(p'')$. In other words, p' is the left-most point of the gap, and p'' is the right-most point of the gap. We know that $\alpha_0 < p' < p'' < 1$ since α_0 and 1 belong to the support of F_ℓ, which is atomless. Notice that F_c must also be constant on $[p', p'')$. Otherwise, any mass placed on that interval by F_c can be moved to an atom at p''. Indeed, this deviation does not change the set of events where the challenger sells (which has positive measure), and only increases the price of sale.

For $\varepsilon > 0$, consider now the alternate pricing strategy for the market leader which equals $F_\ell(p' - \varepsilon)$ for all $q \in [p' - \varepsilon, p'')$, and coincides with F_ℓ elsewhere. There are two opposing effects from switching to this strategy. A loss in profit occurs in comparison to charging $q \in [p' - \varepsilon, p']$, from two sources. There is an increase in the probability that the market draws attention when charging p'' instead of q. There is also an increase in the probability that, if the market is examined, the challenger's price will be lower. Since F_c is constant on $[p', p'')$ and both F_ℓ, F_c are atomless, both increases in probability can be made arbitrarily small by decreasing ε. A gain in profit occurs in comparison to charging $q \in [p' - \varepsilon, p]$, whose magnitude is bounded below by a positive number independent of ε (which comes from selling to fully inattentive consumers at a higher price). Thus this deviation is profitable for small ε, contradicting the existence of a gap in F_ℓ. ∎

The proof of Proposition 7 may give the impression that an analogous argument also applies to F_c. Since we know that F_c can have gaps, where does the logic break down? To explore this, suppose F_c is constant between p' and p'', where $\alpha_0 < p' < p'' < \overline{p}_c$. For simplicity, let us just compare the leader's profit from charging $p' - \varepsilon$ versus p''. When ε is sufficiently small so that $F_c(p' - \varepsilon)$ is close to $F_c(p'')$, the change in profit from switching to p'' is approximately equal to

$$p'' - p' + \varepsilon + F_c(p'')\Big(p'\Pi_\ell\big(F_\ell(p' - \varepsilon)\big) - p''\Pi_\ell\big(F_\ell(p'')\big)\Big).$$

Since both Π_ℓ and F_ℓ are strictly increasing, $p'\Pi_\ell\big(F_\ell(p' - \varepsilon)\big) < p''\Pi_\ell\big(F_\ell(p'')\big)$. The net effect of the deviation could thus be negative even when ε is arbitrarily close to zero. The intuition is that when the leader raises his price from $p' - \varepsilon$ to p'', he faces a strictly higher probability of drawing attention to his market.

This means that there is a higher probability of losing the market when the challenger quotes a price lower than $p' - \varepsilon$, an event whose probability is bounded from zero. Because other leaders do charge prices in (p', p'') with positive probability, this loss can no longer be made arbitrarily small.

The above results allow us to now complete our characterization of the leaders' and challengers' equilibrium pricing strategies. Given that there is zero probability of ties, and given our knowledge of equilibrium profits and firms' highest and lowest prices, the indifference conditions for equilibrium indeed correspond to (4) and (5). Consequently, F_ℓ must be given by

$$F_\ell(p) = \begin{cases} \Pi_c^{-1}\left(\frac{\alpha_0 EA(\alpha)}{Mp}\right) & \text{for all } p \text{ in the support of } F_c, \\ \Pi_\ell^{-1}\left(\frac{p-\alpha_0}{pF_c(p)}\right) & \text{for all other } p \in [\alpha_0, 1], \end{cases}$$

as claimed in Theorem 1. Moreover, for any price in the challengers' support, F_c must coincide with the function $\tilde{F}_c$, which is defined in (7) by

$$\tilde{F}_c(p) = \frac{p - \alpha_0}{p\Pi_\ell\Big(\Pi_c^{-1}\big(\frac{\alpha_0 EA(\alpha)}{Mp}\big)\Big)}.$$

Proving that $F_c(p) = \min_{\tilde{p}\in[p,1]} \tilde{F}_c(\tilde{p})$ for all prices in $[\alpha_0, \bar{p}_c]$, as claimed in Theorem 1, requires one more result.

Proposition 8. *If the price p is in the support of the challengers' strategy, then $\tilde{F}_c(p) \leq \tilde{F}_c(\tilde{p})$ for all $\tilde{p} \in [p, 1]$.*

Proof. This is immediate if $\tilde{p}$ is also in the support of F_c, since in that case $\tilde{F}_c(p) = F_c(p) \leq F_c(\tilde{p}) = \tilde{F}_c(\tilde{p})$, with the inequality following from $p < \tilde{p}$. Suppose then that $\tilde{p}$ is not in the support of F_c and, by contradiction, that $\tilde{F}_c(\tilde{p}) < \tilde{F}_c(p)$. The challenger's profit when charging $\tilde{p}$ is given by

$$\tilde{p}\Pi_c(F_\ell(\tilde{p})) = \tilde{p}\Pi_c\Big(\Pi_\ell^{-1}\big(\frac{\tilde{p} - \alpha_0}{\tilde{p}F_c(\tilde{p})}\big)\Big).$$

Since $p < \tilde{p}$ and p is in the support of the challengers' strategy, $\tilde{F}_c(p) = F_c(p) \leq$

$F_c(\tilde{p})$. Hence $\tilde{F}_c(\tilde{p}) < F_c(\tilde{p})$. Since Π_ℓ is strictly increasing and Π_c is strictly decreasing,

$$\tilde{p}\Pi_c(F_\ell(\tilde{p})) > \tilde{p}\Pi_c\Big(\Pi_\ell^{-1}(\frac{\tilde{p}-\alpha_0}{\tilde{p}\tilde{F}_c(\tilde{p})})\Big).$$

Applying the definition of $\tilde{F}_c$, we conclude that

$$\tilde{p}\Pi_c\Big(\Pi_\ell^{-1}(\frac{\tilde{p}-\alpha_0}{\tilde{p}\tilde{F}_c(\tilde{p})})\Big) = \frac{\alpha_0 EA(\alpha)}{M},$$

which is the challenger's equilibrium profit. Hence F_c could not be part of an equilibrium, since charging $\tilde{p}$ would be a strictly profitable deviation. ∎

The characterization of F_c in Theorem 1 now follows. Indeed, for any price p in the support of F_c, we know that $F_c(p) = \tilde{F}_c(p)$. By Proposition 8, it must be that $\tilde{F}_c(p) \leq \tilde{F}_c(\tilde{p})$ for all $\tilde{p} > p$, proving the desired characterization for those prices that the challenger employs. But the characterization also holds for any price $p < \bar{p}_c$ which is part of a gap in the support of F_c. To see this, observe that the leftmost endpoint of the gap (denoted p_1) and the rightmost endpoint of the gap (denoted p_2) do belong to the support of F_c, and so the desired characterization holds for them. Because F_c is atomless, $F_c(p_1) = F_c(p) = F_c(p_2)$, which squeezes $F_c(p)$ to the desired value.

4.2 Establishing existence

The above results establish uniqueness of equilibrium, if an equilibrium exists. To prove existence, we must argue that F_c and F_ℓ, as described in Theorem 1, are cumulative distribution functions and are, indeed, part of an equilibrium. Let us momentarily defer some of the more technical aspects of the problem, to first persuade the reader that no player would have a profitable deviation.

Remember that Proposition 1 characterizes a consumer's optimal attention allocation. In particular, when the lowest prices of leaders' and challengers' coincide, and there is no chance of ties in leaders' prices, a consumer with k units of attention will optimally examine the markets of the k most expensive leaders. Of course, such a consumer optimally purchases at the lowest price he finds. Since consumers are acting optimally, it remains to show that neither leaders nor challengers have a profitable deviation. The construction of F_c ensures that quoting any price $p \in [\alpha_0, 1]$ gives the leader a profit of α_0. Quoting a price above 1 or a price below α_0 thus yields the leader a strictly smaller profit. The construction of F_ℓ ensures that quoting any price in the support of the challenger's strategy yields a profit of $\alpha_0 EA(\alpha)/M$. Since $EA(\alpha)/M$ is the expected proportion of consumers checking his market, the challengers' profit is clearly larger than that attained by quoting a price smaller than α_0. We now prove that quoting any price $p \geq \alpha_0$ which is not in the support of F_c also yields a smaller profit. Consider any p outside the support of F_c. By construction, $F_c(p) \leq \tilde{F}_c(p)$. Hence

$$F_\ell(p) = \Pi_\ell^{-1}\Big(\frac{p - \alpha_0}{pF_c(p)}\Big) \geq \Pi_\ell^{-1}\Big(\frac{p - \alpha_0}{p\tilde{F}_c(p)}\Big).$$

Applying the decreasing function Π_c on both sides, multiplying by p, and plugging in the definition of $\tilde{F}_c$, we find that $p\Pi_c(F_\ell(p)) \leq \alpha_0 EA(\alpha)/M$. In other words, the challenger cannot obtain a higher profit by quoting a price outside the support of F_c.

Completing the proof of existence requires two additional points, which are provided by Proposition 9 below. First, we must show that F_ℓ and F_c have the properties that allow us to describe firms' profits as we have done above. Second, we must tackle the matter of well-definedness. It not obvious that F_ℓ and F_c are cumulative distribution functions (which requires taking the values 0 and 1 at the appropriate ends, and being increasing). In addition, because the functions Π_ℓ and Π_c do not take all values in $[0, 1]$, we must check that they are being inverted over the appropriate intervals. Notice that the probability Π_ℓ of a market being examined is at least α_M and at most $1 - \alpha_0$ (remember that $\alpha_0 + \alpha_M < 1$); and the probability Π_c that a challenger sells is at least 0 and at most $EA(\alpha)/M$ in a partially symmetric equilibrium.

Proposition 9. *The pricing strategies F_ℓ and F_c are well-defined, atomless cumulative distribution functions. Moreover, F_ℓ is strictly increasing over $[\alpha_0, 1]$, and α_0 is the lowest price in the support of F_ℓ and F_c.*

Proof. We begin with F_c. Observe that $\alpha_0 EA(\alpha)/Mp$ belongs to the range of Π_c for any $p \in [\alpha_0, 1]$, and that the domain of Π_ℓ is $[0,1]$. Hence both $\tilde{F}_c$ and F_c are well-defined at any such p. We argue that F_c is a valid distribution function. It is increasing and continuous by construction. Moreover, it is easy to see that $\tilde{F}_c(\alpha_0) = 0$, as the numerator is zero and the denominator is nonzero: observe that $\Pi_c^{-1}(EA(\alpha)/M) = 0$ and $\Pi_\ell(0) = \alpha_M > 0$. It remains to show that $F_c(\bar{p}_c) = 1$ for some $\bar{p}_c \in (\alpha_0, 1)$, which itself follows if there exists a largest price strictly smaller than one such that $\tilde{F}_c$ equals one. Such a price exists by the Intermediate Value Theorem, because $\tilde{F}_c$ is continuous, with $\tilde{F}_c(\alpha_0) < 1$ and $\tilde{F}_c(1) > 1$. To see the last fact, suppose to the contrary that $\tilde{F}_c(1)$ were less than or equal to one. In that case, we would have $\Pi_c^{-1}(\alpha_0 EA(\alpha)/M) \geq \Pi_\ell^{-1}(1-\alpha_0) = 1$, which is impossible because Π_c is strictly decreasing and satisfies $\Pi_c(1) = 0$. It can be checked by elementary calculus that $\tilde{F}_c{}'(\alpha_0) > 0$, so α_0 is in the support of F_c.

We next show that F_ℓ is well-defined. Again, because $\alpha_0 EA(\alpha)/Mp$ belongs to the range of Π_c for any $p \in [\alpha_0, 1]$, we know that F_ℓ is well-defined whenever p belongs to the support of F_c. Consider then a price $p \in [\alpha_0, 1]$ that does not belong to the support of F_c. Since $F_c(p) \leq \tilde{F}_c(p)$, we have

$$\frac{p-\alpha_0}{pF_c(p)} \geq \frac{p-\alpha_0}{p\tilde{F}_c(p)} = \Pi_\ell\Big(\Pi_c^{-1}\big(\frac{\alpha_0 EA(\alpha)}{Mp}\big)\Big),$$

which is greater than or equal to α_M, as desired. Moreover, we claim that $(p-\alpha_0)/pF_c(p) \leq 1-\alpha_0$. This is obvious if $F_c(p) = 1$. If $F_c(p) < 1$, then there exists some $p' > p$ in the support of F_c such that $F_c(p') = F_c(p)$. Hence

$$\frac{p-\alpha_0}{pF_c(p)} \leq \frac{p'-\alpha_0}{p'F_c(p')} = \Pi_\ell\Big(\Pi_c^{-1}\big(\frac{\alpha_0 EA(\alpha)}{Mp'}\big)\Big),$$

which is less than or equal to $1-\alpha_0$, as desired. Therefore, $(p-\alpha_0)/pF_c(p)$ also belongs to the range of Π_ℓ, ensuring that F_ℓ is also well-defined for prices $p \in [\alpha_0, 1]$ outside of the support of F_c.

Finally, we show that F_ℓ is an atomless and gapless cumulative distribution function. Since α_0 is in the support of F_c, we have $F_\ell(\alpha_0) = 0$. Since 1 is not in the support of F_c, we conclude that $F_\ell(1) = \Pi_\ell^{-1}(1-\alpha_0) = 1$. We complete the proof by showing that F_ℓ as defined in (6) is continuous and strictly increasing over $[\alpha_0, 1]$, which also proves that α_0 is in its support. Since Π_c^{-1} is strictly decreasing and Π_ℓ^{-1} is strictly increasing, each of the two functions defining F_ℓ in (6) is strictly increasing within any interval of prices for which they are applied. Moreover, each of these functions is continuous. The argument is complete if we show that F_ℓ is itself continuous. Let p be a boundary point of the support of F_c, and let $(p_n)_n$ be a sequence which is not in the support of F_c but which converges to p. Since the support of a distribution is closed, p is in the support of F_c and so $F_\ell(p) = \Pi_c^{-1}(\alpha_0 EA(\alpha)/Mp)$. Moreover, because p is in the support of F_c, the minimum in (8) is achieved by $\hat{p} = p$, or $F_c(p) = (p - \alpha_0)/p\Pi_\ell(F_\ell(p))$. Since F_c is continuous, $F_c(p_n)$ converges to $F_c(p)$, and so $F_\ell(p_n)$ converges to $\Pi_\ell^{-1}((p - \alpha_0)/pF_c(p))$. But simple algebra shows $\Pi_\ell^{-1}((p - \alpha_0)/pF_c(p)) = \Pi_c^{-1}(\alpha_0 EA(\alpha)/Mp)$ if and only if $F_c(p) = (p - \alpha_0)/p\Pi_\ell(F_\ell(p))$, completing the proof. ∎

5 Conclusion

This paper proposes a stylized model of price competition, with consumers optimally deciding which components of their expenses to audit given bounds on their attention. In the classic framework, where consumers are fully attentive, the cross-market implications of prices are limited to income and substitution effects. Limited attention brings a new dimension to competition, with the prices of the most visible firms exerting an externality on other markets by deflecting or drawing consumers' attention. Taking into account the firms equilibrium response, decreasing the average attention level benefits consumers through competition for their inattention.

Our model suggests interesting new avenues for exploration. A first direction would be to embed the model into a dynamic framework to determine endogenously which firms serve as default providers. Competition for inattention may be exacerbated, with default providers further lowering their prices, as the benefit of remaining in their position increases the incentive to be under the consumers' radar. A second direction would be to further investigate consumers' optimal allocation of attention in heterogeneous markets. Inspecting markets with the highest expected savings may translate into more intricate attention strategies. A third direction would be to include multiple challengers in each market. Our assumption of a single challenger is a reduced-form representation of friction in identifying challengers and learning their offers. In a more general model, sampling each additional challenger's price would deplete some of the consumer's budget for attention. One can then study the tradeoff between allocating attention across markets versus within markets. A consumer would allocate each additional unit of attention to the market with the highest expected savings given the prices he has observed so far. A fourth direction would be to consider general preferences, allowing for complementarity and non-satiation, to investigate the effect that competition for inattention has on the total surplus.

We hope that the present paper motivates researchers to investigate these questions, and will be useful for further analysis of consumers' optimal allocation of attention and the implications for price theory.

Appendix

Proof of Theorem 2

The result is established in four steps.

Step 1. *If $\tilde{F}_c$ is strictly increasing, then the support of F_c is $[\alpha_0, \bar{p}_c]$ and*

$$F_\ell(p) = \max\Big\{\Pi_c^{-1}\big(\frac{\alpha_0 EA(\alpha)}{Mp}\big)\ ,\ \Pi_\ell^{-1}\big(1-\frac{\alpha_0}{p}\big)\Big\}.$$

Proof. We know $\tilde{F}_c(\alpha_0) < 1 < \tilde{F}_c(1)$ from Proposition 9. Since $\tilde{F}_c$ is strictly increasing and continuous, there is a unique $\bar{p}_c \in (\alpha_0, 1)$ solving $\tilde{F}_c(p_c) = 1$. Using Theorem 1 and increasingness of $\tilde{F}_c$, we know that $F_c(p) = \tilde{F}_c(p)$ for all $p \in [\alpha_0, \bar{p}_c]$, and hence the support of F_c is $[\alpha_0, \bar{p}_c]$. By construction, $F_c(p) < 1$ if and only if $p < \bar{p}_c$. Using the definition of F_c, this means that $\Pi_\ell^{-1}\big(1-\frac{\alpha_0}{p}\big) < \Pi_c^{-1}\big(\frac{\alpha_0 EA(\alpha)}{Mp}\big)$ for $p \in [\alpha_0, \bar{p}_c)$, with the reverse inequality holding for $p \in [\bar{p}_c, 1]$. The construction of F_ℓ in Theorem 1 then implies that F_ℓ is given by the maximum of these two functions. ∎

Step 2. *If $\Pi_c(0) - \Pi_c(x)$ is strictly log-concave with respect to $x \in [0,1]$, except perhaps at a finite number of points, then $\tilde{F}_c$ is strictly increasing for $p \in [\alpha_0, \bar{p}_c]$.*

Proof. We know that $\Pi_c(0) = \frac{EA(\alpha)}{M}$. Subtracting $\Pi_c\big(\Pi_c^{-1}(\frac{\alpha_0 EA(\alpha)}{Mp})\big) = \frac{\alpha_0 EA(\alpha)}{Mp}$ from the previous equation and dividing by $\Pi_\ell\big(\Pi_c^{-1}(\frac{\alpha_0 EA(\alpha)}{Mp})\big)$ yields[8]

$$\frac{\frac{EA(\alpha)}{M}\frac{p-\alpha_0}{p}}{\Pi_\ell\big(\Pi_c^{-1}(\frac{\alpha_0 EA(\alpha)}{Mp})\big)} = \frac{\Pi_c(0) - \Pi_c\big(\Pi_c^{-1}(\frac{\alpha_0 EA(\alpha)}{Mp})\big)}{\Pi_\ell\big(\Pi_c^{-1}(\frac{\alpha_0 EA(\alpha)}{Mp})\big)}.$$

The LHS is a positive constant times $\tilde{F}_c(p)$. Hence $\tilde{F}_c(p)$ is strictly increasing for $p \in [\alpha_0, 1]$ if and only if the RHS is. By assumption, the derivative of $\frac{\Pi_c(0)-\Pi_c(x)}{\Pi_\ell(x)} = 1/(\log(\Pi_c(0) - \Pi_c(x))'$ is strictly positive on $[0,1]$, except perhaps at finitely many points. Continuity of $\frac{\Pi_c(0)-\Pi_c(x)}{\Pi_\ell(x)}$ implies that it is strictly

[8] We thank Xiaosheng Mu for pointing out this identity.

increasing on $[0,1]$. This concludes the proof, using the change of variable $x = \Pi_c^{-1}(\frac{\alpha_0 EA(\alpha)}{Mp})$, which is a strictly increasing function of p. ∎

Step 3. *The following equivalence holds:*

$$\Pi_c(0) - \Pi_c(x) = \frac{1}{M}\sum_{j=0}^{M}\binom{M}{j}x^j(1-x)^{M-j}\sum_{k=0}^{M}\alpha_k \max\Big\{j - M + k, 0\Big\}.$$

Proof. We first recall some standard definitions and identities. The beta function is $B(a,b) = \int_0^1 t^{a-1}(1-t)^{b-1}dt$. The Euler integral of the first kind implies $B(a,b) = \frac{(a-1)!(b-1)!}{(a+b-1)!}$ for integers a, b. The incomplete beta function is $B(x;a,b) = \int_0^x t^{a-1}(1-t)^{b-1}dt$, and the regularized incomplete beta function is $I_x(a,b) = \frac{B(x;a,b)}{B(a,b)}$, which satisfies $I_x(a,b) = \sum_{j=a}^{a+b-1}\binom{a+b-1}{j}x^j(1-x)^{a+b-1-j}$. Next, observe that for each k,

$$\begin{aligned}
\pi_k^c(0) - \pi_k^c(x) &= \int_0^x \sum_{i=0}^{k-1}\binom{M-1}{i}(1-t)^i t^{M-1-i}dt \\
&= \sum_{i=0}^{k-1}\binom{M-1}{i}B(x; M-i, i+1) \\
&= \sum_{i=0}^{k-1}\binom{M-1}{i}B(M-i, i+1)I_x(M-i, i+1) \\
&= \frac{1}{M}\sum_{i=0}^{k-1}I_x(M-i, i+1) \\
&= \frac{1}{M}\sum_{i=0}^{k-1}\sum_{j=M-i}^{M}\binom{M}{j}x^j(1-x)^{M-j} \\
&= \frac{1}{M}\sum_{j=M-k+1}^{M}\Big(j-(M-k)\Big)\binom{M}{j}x^j(1-x)^{M-j},
\end{aligned}$$

since in the penultimate summation, $j = M$ appears k times, $j = M-1$ appears $k-1$ times, $\ldots$, and $j = M-k+1$ appears one time.

Using the above result and interchanging the order of summation,

$$\begin{aligned}
\Pi_c(0) - \Pi_c(x) =& \frac{1}{M}\sum_{k=1}^{M}\alpha_k \sum_{j=M-k+1}^{M}\Big(j-(M-k)\Big)\binom{M}{j}x^j(1-x)^{M-j} \\
=& \frac{1}{M}\sum_{k=1}^{M}\alpha_k \sum_{j=0}^{M}\max\Big\{j-(M-k),0\Big\}\binom{M}{j}x^j(1-x)^{M-j} \\
=& \frac{1}{M}\sum_{j=0}^{M}\binom{M}{j}x^j(1-x)^{M-j}\sum_{k=0}^{M}\alpha_k \max\Big\{j-(M-k),0\Big\} \qquad \blacksquare
\end{aligned}$$

Step 4. *If $(\alpha_1, \ldots, \alpha_M)$ is a log-concave sequence, or if $\alpha_i \leq 2\alpha_j$ for all $i < j$, then $\Pi_c(0) - \Pi_c(x)$ is strictly log-concave in $x \in [0,1]$, excepts perhaps at finitely many points.*

Proof. We will apply the main theorem of Mu (2013), which states that if $(\beta_0, \ldots, \beta_M)$ is a non-constant log-concave sequence, then $\sum_{j=0}^{M}\binom{M}{j}x^{M-j}(1-x)^j\beta_j$ is strictly log-concave in $x \in [0,1]$, except at perhaps finitely many points. Using a change of variable and symmetry of binomial coefficients, observe that $\sum_{j=0}^{M}\binom{M}{j}x^{M-j}(1-x)^j\beta_j = \sum_{j=0}^{M}\binom{M}{j}x^j(1-x)^{M-j}\beta_{M-j}$. If a sequence is log-concave, then it is also log-concave when read backwards. Thus, Mu's theorem holds also when replacing $\sum_{j=0}^{M}\binom{M}{j}x^{M-j}(1-x)^j\beta_j$ by $\sum_{j=0}^{M}\binom{M}{j}x^j(1-x)^{M-j}\beta_j$. Using Step 3, to ensure the desired property of $\Pi_c(0) - \Pi_c(x)$, it thus suffices to show that each of the above properties of $(\alpha_1, \ldots, \alpha_M)$ implies that $(\beta_0, \ldots, \beta_M)$ is log-concave, where we define $\beta_j := \sum_{k=0}^{M}\alpha_k \max\{j-(M-k),0\}$. (Notice that β is non-constant since $\alpha \neq 0$.) Defining $\hat{\alpha}_{M-k} := \alpha_k$, observe that $\beta_j = \sum_{i=0}^{M}\hat{\alpha}_i \max\{j-i,0\}$.

Consider first the case that α is a log-concave sequence (hence so is $\hat{\alpha}$). Since $\max\{i,0\}$ is a log-concave sequence, then so is $\max\{j-i,0\}$. Because each β_j is the convolution of two log-concave sequences, β is log concave itself. Next, consider the case that $\alpha_i \leq 2\alpha_j$ when $i < j$. Applying the identity

$\beta_k = \beta_{k-1} + \sum_{i=1}^{k} \alpha_{M-i+1}$, and rearranging terms, β is log-concave iff

$$\begin{aligned}
& \beta_k(\beta_{k-1} + \sum_{i=1}^{k} \alpha_{M-i+1}) \geq \beta_{k-1}\beta_{k+1} \\
\Leftrightarrow \quad & \beta_k \sum_{i=1}^{k} \alpha_{M-i+1} \geq \beta_{k-1} \sum_{i=1}^{k+1} \alpha_{M-i+1} \\
\Leftrightarrow \quad & \Big(\sum_{i=1}^{k} \alpha_{M-i+1}\Big)^2 \geq \beta_{k-1}\alpha_{M-k} \\
\Leftrightarrow \quad & \sum_{i=1}^{k} \alpha_{M-i+1}^2 + 2\sum_{i=1}^{k-1} \sum_{j=i+1}^{k} \alpha_{M-i+1}\alpha_{M-j+1} \geq \beta_{k-1}\alpha_{M-k}.
\end{aligned}$$

Note that $2\alpha_{M-j+1} \geq \alpha_{M-k}$ when $j \leq k$. Hence the left-hand side of the last expression is at least $\alpha_{M-k} \sum_{i=1}^{k-1}(k-i)\alpha_{M-i+1}$ which is precisely $\alpha_{M-k}\beta_{k-1}$, as desired. This concludes the proof of this step and of Theorem 2 ∎

Proof of Theorem 4

Given the discussion following Theorem 4 in the text, all that remains to show is that $\hat{\Pi}_c(x)/\Pi_c(x)$ is decreasing in x, and strictly so at $x = 0$.

It will be convenient to prove this in a more general setting, where the distributions α and $\hat{\alpha}$ come from a family of distributions parametrized by λ, a real-valued scalar taking values in either a continuous or discrete set; the case where λ can take one of two values corresponds to the presentation in Section 3.3. Let $\alpha(\lambda)$ denote the distribution from this family given λ. The family satisfies (i) $\alpha_0(\lambda) = \alpha_0(\hat{\lambda})$ for all $\hat{\lambda} \neq \lambda$; (ii) the partial attention distribution $a(\alpha(\lambda))$ is a log-concave sequence for each λ; and (iii) the monotone likelihood ratio property on the partial attention distribution $a(\alpha(\cdot))$. Property (iii) is equivalent to $\frac{\alpha_k(\hat{\lambda})}{\alpha_k(\lambda)} \leq \frac{\alpha_{k+1}(\hat{\lambda})}{\alpha_{k+1}(\lambda)}$ for $\hat{\lambda} > \lambda$ and $k \in \{1, \ldots, M-1\}$.

For any λ, let $\Pi_\ell(\lambda, x) = \sum_{k=1}^{M} \alpha(\lambda)\pi_k^\ell(x)$ and $\Pi_c(\lambda, x) = \sum_{k=1}^{M} \alpha(\lambda)\pi_k^c(x)$. Also, let $EA(\alpha(\lambda))$ be the expected level of attention under $\alpha(\lambda)$. Note that the MLR ranking on the partial attention distributions amounts to log-supermodularity in k, λ. Similarly, note that for $\hat{\lambda} > \lambda$, decreasingness of

$\Pi_c(\hat{\lambda}, x)/\Pi_c(\lambda, x)$ in x amounts to log-submodularity of $\Pi_c(\lambda, x)$ in λ, x. The proof then proceeds in two steps.

Step 1. $\Pi_c(\lambda, x)$ *is log-submodular in* λ, x.

Proof. It is well-known that if the function $t(i, y)$ is log-supermodular in i, y and the function $s(i, z)$ is log-supermodular in i, z, then $\int_i t(i, y)s(i, z)di$ is log-supermodular in y, z (see, for example, Corollary 1 in Quah and Strulovici, 2011). This preservation of log-supermodularity result extends to discrete summations (e.g., i comes from the set $\{1, 2, \ldots, n\}$).[9] To see this, apply the preservation result to the functions $\tilde{t}(j, y)$ and $\tilde{s}(j, z)$, which are defined with $j \in [0, 1)$ as follows: if $\frac{i-1}{n} \leq j < \frac{i}{n}$, then $\tilde{t}(j, y) = t(i, y)$ and $\tilde{s}(j, y) = s(i, y)$. Below, we iteratively apply the preservation result to prove that $\Pi_c(\lambda, x)$ is log-submodular in λ, x. Consider the function

$$\int_0^1 1_{(t \leq x)} \left(\sum_{k=1}^{M} \alpha_k(\lambda) \sum_{i=1}^{M} 1_{(i \leq k-1)} \binom{M-1}{i} t^i (1-t)^{M-i-1} \right) dt, \tag{10}$$

which is simply $\Pi_c(\lambda, 1-x)$ using a change of variables from t to $1-t$ (note that $1(\cdot)$ is the indicator function which is equal to 1 if its argument is true). We first show that $1_{i \leq k-1}$ is log-supermodular in i and k. Indeed, consider $(\bar{i}, \bar{k}) \geq (\underline{i}, \underline{k})$. Then

$$1_{(\bar{i} \leq \bar{k}-1)} 1_{(\underline{i} \leq \underline{k}-1)} \geq 1_{(\bar{i} \leq \underline{k}-1)} 1_{(\underline{i} \leq \bar{k}-1)},$$

since if the right-hand side is one, then so is the left-hand side. Next, we show that $\binom{M-1}{i} t^i (1-t)^{M-i-1}$ is log-supermodular in i, t. Indeed, the ratio

$$\frac{\binom{M-1}{i} t^i (1-t)^{M-i-1}}{\binom{M-1}{i-1} t^{i-1} (1-t)^{M-i}} = \frac{(M-i)t}{i(1-t)}$$

is increasing in t for $t \in [0, 1)$. Applying the preservation result, this implies that the inner sum in (10) is log-supermodular in k, t. By assumption, $\alpha_k(\lambda)$

[9] We thank Bruno Strulovici for pointing this out.

is log-supermodular in k, λ. Applying the preservation result again, the expression inside the large parentheses in (10) is log-supermodular in t, λ. Since $1_{(t \leq x)}$ is log-supermodular in t, x (the argument is the same as before), the entire expression in (10) is log-supermodular in λ, x, by applying the standard preservation result. But since that sum is $\Pi_c(\lambda, 1-x)$, we obtain that $\Pi_c(\lambda, x)$ is log-submodular in λ, x as desired. ∎

Step 2. *The derivative of* $\Pi_c(\hat{\lambda}, x)/\Pi_c(\lambda, x)$ *with respect to* x *is strictly negative at* $x = 0$ *for* $\hat{\lambda} > \lambda$.

Proof. The sign of this derivative is the same as the sign of

$$\Pi_c(\hat{\lambda}, 0)\Pi_\ell(\lambda, 0) - \Pi_c(\lambda, 0)\Pi_\ell(\hat{\lambda}, 0) = \frac{EA(\alpha(\hat{\lambda}))}{M}\alpha_M(\lambda) - \frac{EA(\alpha(\lambda))}{M}\alpha_M(\hat{\lambda}).$$

This expression is proportional to

$$\alpha_M(\lambda)\sum_{k=1}^{M} k\alpha_k(\hat{\lambda}) - \alpha_M(\hat{\lambda})\sum_{k=1}^{M} k\alpha_k(\lambda) = \sum_{k=1}^{M-1} k[\alpha_M(\lambda)\alpha_k(\hat{\lambda}) - \alpha_M(\hat{\lambda})\alpha_k(\lambda)],$$

which is indeed strictly negative because $\frac{\alpha_k(\hat{\lambda})}{\alpha_k(\lambda)} \leq \frac{\alpha_M(\hat{\lambda})}{\alpha_M(\lambda)}$ for all k by log-supermodularity, with at least one strict inequality. ∎

References

Armstrong, Mark and Yongmin Chen (2009), "Inattentive Consumers and Product Quality", *Journal of the European Economic Association*, 7, 411–422.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer (2013), "Competition for Attention", *Working Paper.*

Burdett, Kenneth and Kenneth Judd (1983), "Equilibrium Price Dispersion", *Econometrica*, 51(4), 955–969.

Dasgupta, Partha and Eric Maskin (1986), "The Existence of Equilibrium in Discontinuous Economic Games, I: Theory", *Review of Economic Studies*, 53(1), 1-26.

Eliaz, Kfir and Ran Spiegler (2011a), "Consideration Sets and Competitive Marketing", *Review of Economic Studies*, 78(1), 235–262.

Eliaz, Kfir and Ran Spiegler (2011b), "On the Strategic Use of Attention Grabbers", *Theoretical Economics*, 6(1), 127–155.

Gabaix, Xavier and David Laibson (2006), "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets", *Quarterly Journal of Economics*, 121(2), 505–540.

Milgrom, Paul (1981), "Good News and Bad News: Representation Theorems and Applications", *The Bell Journal of Economics*, 12(2), 380–391.

Mu, Xiaosheng (2013), "Log-concavity of a Convex Combination of Binomial Distribution Functions", *Working Paper.* Available for download at https://sites.google.com/site/xiaoshengswebpage/

Quah, John and Bruno Strulovici (2012), "Aggregating the Single Crossing Property", *Econometrica*, 80(5), 2333–2348.

Reny, Philip J. (1999), "On the Existence of Pure and Mixed Strategy Nash Equilibria in Discontinuous Games", *Econometrica*, 67(5), 1029–1056.

Rosenthal, Robert W. (1980), "A Model in which an Increase in the Number of Sellers Leads to a Higher Price", *Econometrica*, 48(6), 1575–1579.

Salop, Steven and Joseph Stiglitz (1977), "Bargains and Ripoffs: A Model of Monopolistically Competitive Price Dispersion", *Review of Economic Studies,* 44(3), 493–510.

Sims, Christopher (2003), "Implications of Rational Inattention", *Journal of Monetary Economics*, 50, 665–690.

Spiegler, Ran (2006), "Competition over Agents with Boundedly Rational Expectations", *Theoretical Economics*, 1(2), 207–231.

Spiegler, Ran (2011), *Bounded Rationality and Industrial Organization*, Oxford University Press, USA.

Stahl, Dale (1989), "Oligopolistic Pricing with Sequential Consumer Search", *American Economic Review*, 79, 700–712.

Varian, Hal R. (1980), "A Model of Sales", *American Economic Review*, 70(4), 651–659.

Woodford, Michael (2009), "Information-Constrained State-Dependent Pricing", *Journal of Monetary Economics*, 56, Supplement S100–S659.