

This item was submitted to [Loughborough's Research Repository](#) by the author.
Items in Figshare are protected by copyright, with all rights reserved, unless otherwise indicated.

Sales competition with heterogeneous firms

PLEASE CITE THE PUBLISHED VERSION

<http://www.eea-esem-congresses.org/index.php?sezn=0&page=77>

PUBLISHER

European Economic Association

VERSION

AM (Accepted Manuscript)

PUBLISHER STATEMENT

This work is made available according to the conditions of the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International (CC BY-NC-ND 4.0) licence. Full details of this licence are available at:
<https://creativecommons.org/licenses/by-nc-nd/4.0/>

LICENCE

CC BY-NC-ND 4.0

REPOSITORY RECORD

Shelegia, Sandro, and Chris M. Wilson. 2019. "Sales Competition with Heterogeneous Firms". figshare.
<https://hdl.handle.net/2134/22246>.

Sales Competition with Heterogeneous Firms

Sandro Shelegia and Chris M. Wilson*

15 February 2016

Abstract

To better understand sales behavior and price dispersion, this paper presents a substantially generalized clearinghouse framework. The framework can permit multiple dimensions of firm heterogeneity, and views firms as competing directly in utility rather than prices. As such, the paper can i) reproduce and extend many equilibria from the existing literature, ii) offer a range of new results on how firm heterogeneity affects sales behavior and market performance, iii) provide original insights into the number and type of firms that engage in sales, and iv) offer a foundation to assess and extend current empirical procedures.

Keywords: Sales; Price Dispersion; Advertising; Clearinghouse; Utility Space

JEL Codes: L13; D43; M3

*Shelegia: Department of Economics and Business, Universitat Pompeu Fabra, sandro.shelegia@upf.edu. Wilson: School of Business and Economics, Loughborough University, c.m.wilson@lboro.ac.uk. We thank Mark Armstrong, Mike Baye, Yossi Spiegel, and seminar participants at the Workshop on Consumer Search and Switching Costs (Indiana) the Workshop on the Economics of Advertising and Marketing (Vienna), Industrial Organization: Theory, Empirics and Experiments Conference (Alberobello), the University of Vienna, Universitat Autònoma de Barcelona, and Brunel University.

1 Introduction

The evidence for the existence of price dispersion within markets is overwhelming, even when products are homogeneous (as reviewed by Baye et al. (2006)). While some of this dispersion arises from persistent differences between firms, such as service levels, empirical findings suggest that much of it arises from temporary price reductions or ‘sales’. Indeed, such sales activity is estimated to account for 20-50% of retail price variation and 38% of all packaged consumer good purchases in the US.¹ Consequently, sales are an active and important research area in many disciplines.²

One major stream of the theoretical literature documents how sales arise in equilibrium due to variation in consumers’ search frictions, or due to the existence of moderate advertising costs.³ However, due to the associated technical complexities, existing models only consider markets with symmetric firms, or markets with weak forms of firm heterogeneity under restrictive assumptions.

This is a significant limitation for several reasons. First, firm heterogeneity is prevalent in practice across many dimensions including products, production costs, service levels, customer bases, and advertising costs. As detailed later in the paper, such heterogeneity is consistent with typical empirical findings that show that firms employ sales, but differ in their average price levels. Second, the limited analysis of firm heterogeneity restricts our understanding of sales competition regarding issues such as how firm characteristics affect sales and market performance, or how agents may wish to influence the degree of market asymmetry. Among other settings, these issues are key when addressing policy concerns such as mergers or vertical relations in sales markets. Third, the focus on symmetric models also limits empirical research. Without a general foundation, current empirical papers are forced to use a restrictive approach that attempts to ‘clean’ the data from the effects of firm heterogeneities before applying the insights of a symmetric sales model.

To help resolve these limitations, this paper presents a flexible and tractable framework that can simultaneously permit multiple dimensions of firm heterogeneity. The framework can i) reproduce and extend many equilibria from the existing literature, ii) offer a range of new results on how firm heterogeneity affects sales behavior and market performance, iii) provide original insights into the number and type of firms that engage in sales behavior, and iv) offer a deeper foundation to assess and extend current empirical procedures.

¹See Nakamura and Steinsson (2008), Hosken and Reiffen (2004); and Steenkamp et al. (2005).

²For related reviews, see Baye et al. (2006) and Raju (1995).

³For example, Varian (1980), Burdett and Judd (1983), Stahl (1989), Janssen and Moraga-González (2004); Robert and Stahl (1993), Baye and Morgan (2001).

In more detail, the paper provides an extended and fully asymmetric version of the ‘clearinghouse’ sales framework (e.g. Baye and Morgan (2001), Baye et al. (2004a), Baye et al. (2006)). In the original model, the market is symmetric and each firm sells a homogeneous product. Firms choose their price and whether to inform consumers of this price given some advertising cost. Consumers are potentially split into ‘loyals’ that only buy from a designated firm, and ‘shoppers’ that exhibit no such loyalty. In equilibrium, as consistent with sales behavior, each firm randomizes between selecting a high price without advertising, and advertising a lower price drawn from a common support.

We modify this model in two important respects. First, we allow firms to compete in utility rather than prices. By drawing on the seminal (non-sales) model of competition in the utility space by Armstrong and Vickers (2001), firms compete directly in utility, u , together with an associated profit function, $\pi(u)$, that captures the maximum profit a firm can make per consumer for a given utility offer. With little increase in computation, this approach provides a high level of generality across many demand, product, cost, and pricing conditions.

Second, we make a subtle change to the consumers’ tie-break rule in a way that can generate enough additional tractability to permit multiple forms of firm heterogeneity. In contrast to the existing literature which uses a tie-break rule that favors advertising firms (e.g. Baye and Morgan (2001), Baye et al. (2004a), Baye et al. (2006) and Arnold et al. (2011)), we consider a tie-break rule where consumers mix between any tied firms with equilibrium probabilities. This distinction makes no difference in symmetric settings. However, it offers additional flexibility under firm heterogeneity by allowing each firm to have the same incentive to advertise a common upper utility bound. This ensures that the resulting equilibrium is tractable while allowing firms to still vary in their advertising probabilities and utility distributions.⁴

In Sections 2 and 3, we present our framework. To ease exposition, we first consider the case of duopoly. After deriving the equilibrium, we demonstrate how the framework can reproduce many equilibria from the existing literature and extend them to allow for more complex market conditions including multiple products, downward-sloping demand,

⁴Specifically, the previous literature assumes that if shoppers expect a non-advertising firm to set some price p , while another firm advertises the same price, then all shoppers trade with the advertising firm. This ‘pro-advertiser’ rule is justified in the listing interpretation of the clearinghouse model (Baye and Morgan (2001)) where any shopper who visits the clearinghouse strictly prefers to buy from a listed firm rather than an equivalent non-listed firm to avoid an additional visit cost. However, we focus more on the advertising interpretation of the clearinghouse model (Baye et al. (2004a) and Baye et al. (2006)) where shoppers receive all adverts before making any visit decisions. Here, the ‘pro-advertiser’ rule seems less reasonable. Instead, as consistent with our tie-break rule, shoppers should be willing to visit any firm with the same expected price.

and positive advertising costs. Among many others, these include symmetric models such as Varian (1980), Baye et al. (2004a), Baye et al. (2006), and Simester (1997), as well as specific asymmetric models, such as Narasimhan (1988), Baye et al. (1992), and Kocas and Kiyak (2006).

In Section 4, we perform a range of comparative statics to offer a number of new insights into how firm heterogeneity affects sales behavior and market performance. In particular, we consider the effects of a change in a single firm’s share of loyal consumers, advertising costs, or profit function. In many cases, these comparative statics differ to the well-known results for symmetric markets.

Section 5 further illustrates the ability of the framework to produce new insights by considering markets with more than two firms. Here, existing research on sales with heterogeneous firms is particularly scant due to the additional complexities involved. However, in a setting with heterogeneous firms, unit demand, and zero advertising costs, the existing literature suggests that only two firms can ever engage in sales behavior while any remaining firms simply price highly to their loyal consumers (Baye et al. (1992), Kocas and Kiyak (2006) and Shelegia (2012)). In contrast, we show that this two-firm result is a special case of our more general findings. In particular, once we allow for positive advertising costs, any number of firms $k^* \in [2, n]$ can engage in equilibrium sales and advertising behavior. While this result appears surprising, its explanation is intuitive. Higher advertising costs soften competition for the shoppers and prompt some relatively inefficient or low quality firms to start using sales despite the increased costs of doing so. This relationship offers an interesting empirical prediction that is yet to be tested. In addition, we then provide a broad characterization of the types of firms that are likely to advertise in equilibrium. For instance, *ceteris paribus*, sales will be used by more profitable firms with relatively low loyal shares and low advertising costs.

Finally, Section 6 details how our framework can be used to assess and extend current empirical procedures. Typical empirical methods attempt to ‘clean’ raw price data from any firm-level heterogeneities by first estimating a price regression with firm fixed-effects, before then using the price residuals as the basis for testing theoretical predictions and/or offering structural estimates of market parameters in sales markets.⁵ However, this approach is known to be restrictive. In the only theoretical justification within the literature, Wildenbeest (2011) shows that the procedure is valid in a setting of unit demand and zero

⁵Examples include Sorensen (2000), Lach (2002), Brown and Goolsbee (2002), Barron et al. (2004), Lewis (2008), Chandra and Tappata (2011), Pennerstorfer et al. (2015), and Sherman and Weiss (2015); Wildenbeest (2011), Moraga-González et al. (2013), Giuliatti et al. (2014), Allen et al. (2014), and An et al. (2015).

advertising costs if firms differ in quality and costs under a specific condition. We exploit our general framework to further explore the procedure in two ways. First, we show that a related specific condition can exist under downward-sloping demand, but that the required empirical procedure becomes much more complex and data-intensive. Second, we suggest a basis for a modified procedure that can consider unit demand settings under a much broader range of firm heterogeneity than that considered by Wildenbeest (2011). These insights should help future empirical work to better understand sales behavior and price dispersion.

Related Literature: Armstrong and Vickers (2001) introduced the concept of competition in the utility space and embedded this approach into a range of discrete-choice pure-strategy equilibrium settings. In contrast, we transfer their utility approach into a qualitatively different asymmetric (clearinghouse) model with advertising, where i) consumers are initially uninformed about firm’s utility offers, and ii) consumers have identical preferences, such that the equilibrium necessarily involves mixing in utility for most parameter values. Some past sales papers have also made reference to competition in utility or ‘surplus’.⁶ However, unlike our framework, these papers only use utility as a means of computing sales equilibria in very specific market settings, and do not use the associated profit function, $\pi(u)$, to explore any general results, implications or wider questions.

As detailed later in the paper, the few existing models of sales with firm heterogeneity often assume single products, unit demand, and zero advertising costs (e.g. Narasimhan (1988), Baye et al. (1992), and Kocas and Kiyak (2006)). Our framework can often reproduce and extend such equilibria. However, our framework cannot reproduce the equilibria in two recent sales papers. First, under unit demand, Arnold et al. (2011) present a clearinghouse duopoly model with positive advertising costs where firms differ in their shares of loyal consumers. However, contrary to our framework, they maintain the past literature’s ‘pro-advertiser’ tie-break rule. This generates a different form of sales equilibrium which, unlike ours, does not converge to standard asymmetric equilibria within the literature when advertising costs tend to zero (e.g. Narasimhan (1988)). Second, Anderson et al. (2015) allow for firm heterogeneity in a non-clearinghouse model without loyal consumers where firms must advertise to earn positive profits. Unlike our model, they find that only two firms can ever engage in sales and advertising behavior when firms are heterogeneous, regardless of the level of advertising costs. As such, they cannot connect to the larger clearinghouse literature or analyze how market factors affect the

⁶See Simester (1997), Hosken and Reiffen (2007), Wildenbeest (2011), Dubovik and Janssen (2012), Anderson et al. (2015).

number and type of firms that use sales. Instead, they focus on some different issues surrounding equilibrium selection and welfare.

2 Model

Let there be two firms, $i = a, b$, and a unit mass of risk-neutral consumers with a zero outside option. Suppose that firm i competes by choosing a utility offer (net of any associated payments), $u_i \in \mathfrak{R}$. Consumers have identical preferences in the sense that all consumers value firm i 's offering at exactly u_i . The maximum possible profit that firm i can extract per consumer when providing u_i is then defined as $\pi_i(u_i)$.

The source of utility and the associated profit function for each firm can be allowed to depend upon a rich set of factors, including demand, product, and cost conditions. However, to provide a simple illustrative example, let firm i sell a single good at price p_i with marginal cost c_i , and suppose consumers have a unit demand with a willingness to pay of V_i . If firm i sets a price p_i , its utility offer then equals $u_i = V_i - p_i$, while its profits per consumer are $\pi_i(u_i) = V_i - c_i - u_i$ for $u_i \geq 0$, and $\pi_i(u_i) = 0$ otherwise.

For each firm, we assume that $\pi_i(u_i)$ is independent of the number of consumers served. We also assume that $\pi_i(u_i)$ is continuously differentiable and strictly quasi-concave in u_i with a unique maximizer at firm i 's monopoly utility level, $u_i^m \geq 0$.⁷ This implies the following. First, as a necessary condition for sales behavior, we require each firm's per consumer monopoly profits to be strictly positive, $\pi_i(u_i^m) > 0$; otherwise firms would only ever be willing to select $u_i = u_i^m$. Second, as $\pi_i(u_i)$ is strictly decreasing in $u_i > u_i^m$, we denote \hat{u}_i as the unique maximum utility offer that firm i can make while breaking even, such that $\pi_i(u_i) < 0$ for all $u_i > \hat{u}_i$. It then follows that $\pi_i(u_i^m) > \pi_i(\hat{u}_i) = 0$ and $\hat{u}_i > u_i^m \geq 0$.

Consumers are initially uninformed about firms' utility offers. Each firm can choose whether or not to inform consumers of its offer through informative advertising under the following assumptions. First, as consistent with all previous clearinghouse models, any advert is observed by all consumers. Second, to maintain our focus, we follow the 'advertising' version of the clearinghouse model (Baye et al. (2004a), Baye et al. (2006)) by keeping advertising costs exogenous. However, in contrast to the past literature, we allow the advertising costs to differ across firms, A_i . Third, each advertising cost is

⁷For ii) note that consumer's optimal behavior is embedded in $\pi_i(u_i)$. Thus, for example, in the unit demand example above, $\pi_i(u_i) = 0$ for $u_i < 0$, and so profits are maximized at $u_i^m = 0$.

strictly positive, with $A_i > 0$ for $i = a, b$.⁸ Now, by defining firm i 's advertising decision as $a_i \in \{0, 1\}$, and given the utility offer and advertising decision of firm $j \neq i$, we can then describe the total proportion of consumers that buy from firm i as $q(u_i, u_j; a_i, a_j)$, together with firm i 's total profits, $\Pi_i = \pi_i(u_i)q(u_i, u_j; a_i, a_j) - a_i A_i$.

Regardless of whether or not firm i has advertised, consumers must visit firm i to make a purchase. Thus, advertising only acts to ensure that consumers know a firm's utility offer before visiting. While these assumptions can be relaxed, we simplify exposition by assuming that consumers can only visit one firm and that the cost of making such a visit is negligible. Therefore, consumers may either visit an advertising firm to buy from its known utility offer, visit a non-advertising firm to discover its utility offer and potentially buy, or immediately exit the market.⁹

Consumers are potentially decomposed into two types, $t \in \{L, S\}$, with proportions, θ and $(1 - \theta)$ respectively, where $\theta \geq 0$. Regardless of any advertised information, 'loyal' consumers, $t = L$, only ever visit their designated 'local' firm. We allow firms to differ in the size of their loyal consumer base, $\theta_i \geq 0$, with $\theta_a + \theta_b = \theta$. The remaining 'shopper' consumers, $t = S$, have no such loyalty. They compare any advertised offers with the utility they expect at any non-advertising firm, and visit the firm with the highest expected utility offer. Having visited a firm, any consumer then buys according to its underlying demand function if the firm offers non-negative utility, and exits otherwise.

We analyze the following game. In Stage 1, each firm simultaneously chooses its utility offer, $u_i \in \mathfrak{R}$, and its advertising decision, $a_i \in \{0, 1\}$. To allow for mixed strategies, define i) $\alpha_i \in [0, 1]$ as firm i 's advertising probability, ii) $F_i^A(u)$ as firm i 's utility distribution when advertising, on support $[\underline{u}^A, \bar{u}^A]$, and iii) $F_i^N(u)$ as firm i 's utility distribution when not advertising, on support $[\underline{u}^N, \bar{u}^N]$. In Stage 2, consumers observe any advertisements and then make their visit and purchase decisions in accordance with the strategies outlined above.

We consider equilibria where all players hold correct beliefs, and where, given all other players' strategies, firms select their utility and advertising strategy optimally, and shoppers have no incentive to change their visit and purchase decisions. Specifically, we

⁸This ensures that our later tie-breaking probabilities are well-defined. However, as later verified, our equilibrium is well-behaved in the sense that when $A_i = A_j \rightarrow 0$, it converges to the equilibrium of a parallel model that allows for $A_i = A_j = 0$ explicitly.

⁹These assumptions can be substantially generalized by allowing consumers to visit the firms sequentially provided that i) the costs of any first visit are not too large, and ii) each consumer may only purchase from a single firm. The latter 'one-stop shopping' assumption may be reasonable in markets such as supermarkets and restaurants, and is frequently assumed in consumer search models, and the wider literature on price discrimination. For technical details, see later Appendix C.

focus on equilibria that involve the following tie-breaking rule: whenever shoppers are indifferent between the offers from advertising firm(s) and/or the expected offers from non-advertising firm(s), shoppers visit firm i with probability x_i , where x_i is determined as part of equilibrium and where $x_a + x_b = 1$. As detailed in the introduction, while natural, this rule differs from the previous literature's 'pro-advertiser' rule, where in the event of a tie between an advertising firm and a non-advertising firm, all shoppers are required to visit the advertising firm (Baye and Morgan (2001), Baye et al. (2004a), Baye et al. (2006), and Arnold et al. (2011)).

While our tie-breaking rule provides a lot of tractability, it is not sufficient to avoid an implicit assumption that is pervasive in the existing literature. We are the first to formally state it:

$$u_a^m = u_b^m = u^m \quad (\text{Assumption U})$$

Assumption U requires that both firms offer the same monopoly utility. However, it makes no restrictions on each firms' monopoly profits, $\pi_i(u^m) \equiv \pi_i^m$. Hence, it actually permits a wide range of settings that go well beyond the previous literature including i) all symmetric settings, ii) any asymmetric setting where consumers have unit demands or where firms use two-part tariffs as then $u_a^m = u_b^m = 0$, and iii) any setting with asymmetric loyal shares and/or advertising costs provided the firms share a common profit function, $\pi_a(u) = \pi_b(u) = \pi(u)$. Outside Assumption U, the tractability provided by our tie-breaking rule is lost: shoppers strictly prefer one firm when neither advertises and any resulting mixed strategy equilibria loses significant elegance, although some qualitative features of our equilibrium remain.

Finally, we assume that both firms have some potential incentive to advertise. Assumption A makes a minimal restriction to ensure that firm i 's profits from not advertising with $u_i = u^m$ and selling only to its loyal consumers, $\theta_i \pi_i^m$, are less than the profits it could obtain by advertising an offer of just above u^m to gain all the shoppers, $(1 - \theta_j) \pi_i^m - A_i$. This assumption is relatively innocuous but allows us to focus on equilibria where both firms advertise.

$$A_i \leq (1 - \theta) \pi_i^m \quad \forall i = a, b \quad (\text{Assumption A})$$

3 Equilibrium Analysis

We proceed in a series of steps. First, any firm that does not advertise will optimally set the monopoly utility level, u^m , with probability one, such that $\underline{u}_i^N = \bar{u}_i^N = u^m$ for $i = a, b$. This follows as i) each firm has monopoly power over its loyal consumers, and ii) if any shoppers search a firm with expectations at or above u^m , then that firm would prefer to offer u^m as consumers cannot purchase elsewhere.

Next, under our tie-breaking rule, any advertised offer of u^m is strictly dominated by not advertising because advertising has no impact on consumer behavior in a tie and yet costs, $A_i > 0$. Hence, firm i will only ever advertise with $u_i > u^m$ such that the lower bound of its advertised utility support is strictly greater than its upper non-advertised utility bound, $\underline{u}_i^A > \bar{u}_i^N = u^m$ for $i = a, b$. Therefore, from this point forward, we can simply refer to firm i 's utility distribution *unconditional on advertising*, $F_i(u)$, with associated support, $[u^m, \bar{u}_i]$. Firm i then advertises with the probability that it offers a utility greater than u^m , $\alpha_i = 1 - F_i(u^m)$.

By selecting u^m and not advertising, firm i will only ever possibly trade with the shoppers if i) firm j also chooses not to advertise, and ii) the shoppers visit i rather than j in the subsequent tie. Consequently, when combined with the revenues from firm i 's loyal consumers, firm i can always guarantee the following total profits by not advertising, for any given x_i and α_j :

$$\pi_i^m[\theta_i + x_i(1 - \alpha_j)(1 - \theta)] \quad (1)$$

Given positive advertising costs, there can be no equilibrium where both firms advertise with probability one because a firm would deviate to avoid the advertising cost. However, as now formalized below, the unique symmetric equilibrium will take one of two forms depending upon the level of advertising costs: when advertising costs are sufficiently large, neither firm advertises, but for lower advertising costs, both firms advertise with interior probabilities.

We first consider the latter form of equilibrium. Suppose, as later derived, that both firms advertise with probabilities, $\alpha_i \in (0, 1)$ for $i = a, b$, and that the tie-breaking probabilities equal x_a and x_b . Then, by adapting standard arguments, one can show that no equilibrium exists with pure utility strategies. Instead, it follows that:

Lemma 1. *In a mixed strategy equilibrium, whenever a firm advertises, it randomizes its utility offer from a common interval $(u^m, \bar{u}]$ without gaps or point masses.*

For each firm to advertise with an interior probability, $\alpha_i \in (0, 1)$ for $i = a, b$, firm i 's profits from not advertising in (1) must equal its profits from advertising an offer slightly higher than u^m , where for a cost of A_i it can win the shoppers outright with the probability that its rival does not advertise, $1 - \alpha_j$. Hence, for both firms, we require $\pi_i^m[\theta_i + x_i(1 - \alpha_j)(1 - \theta)] = \pi_i^m[\theta_i + (1 - \alpha_j)(1 - \theta)] - A_i$. From this, for a given $x_j = 1 - x_i$, we can state:

$$\alpha_i = 1 - \frac{A_j}{x_i(1 - \theta)\pi_j^m}. \quad (2)$$

Now note that each firm must always expect to earn its equilibrium profits, $\bar{\pi}_i$, when mixing over advertising. By substituting α_i from (2) into the profits from not advertising (1), and using $x_j = 1 - x_i$, we obtain (3). Equilibrium profits therefore derive from two channels: the existence of loyal consumers and costly advertising.

$$\bar{\pi}_i = \theta_i \pi_i^m + \frac{x_i}{1 - x_i} A_i \quad (3)$$

Next, we derive the utility distributions. By advertising any offer within $(u^m, \bar{u}]$, firm i gains expected profits of $\pi_i(u)[\theta_i + (1 - \theta)F_j(u)] - A_i$. Intuitively, firm i always collects its loyal consumer profits but also wins the profits of the $(1 - \theta)$ shoppers with the probability that firm j does not advertise a higher utility, $F_j(u)$. For each firm to be indifferent over the advertising support, their respective expected profits must be equal for each $u \in (u^m, \bar{u}]$. Moreover, for each firm to be indifferent between advertising and not advertising, these expected profits must equal the firm's respective equilibrium profits. As there are no mass points within $(u^m, \bar{u}]$, this requires $\bar{\pi}_i = \pi_i(u)[\theta_i + (1 - \theta)F_j(u)] - A_i$ for $i = a, b$. By substituting from (3) and reversing subscripts, firm i 's utility distribution can be expressed by (4), where $F_i(u^m) = 1 - \alpha_i$ as expected.

$$F_i(u) = \frac{\bar{\pi}_j - \theta_j \pi_j(u) + A_j}{(1 - \theta)\pi_j(u)} = \frac{x_i \theta_j [\pi_j^m - \pi_j(u)] + A_j}{x_i(1 - \theta)\pi_j(u)} \quad (4)$$

We now find the upper bound, \bar{u} , and the equilibrium tie-breaking probabilities, x_i and x_j . As there is no point mass at \bar{u} , a firm that advertises \bar{u} will definitely win all the shoppers, such that $\bar{\pi}_i = (1 - \theta_j)\pi_i(\bar{u}) - A_i$. By substituting from (3), we can express x_i with the first equality in (5) below. The second equality then follows by using $1 - x_i = x_j$ and reversing all the subscripts.

$$x_i = 1 - \frac{A_i}{\pi_i(\bar{u})(1 - \theta_j) - \theta_i \pi_i^m} = \frac{A_j}{\pi_j(\bar{u})(1 - \theta_i) - \theta_j \pi_j^m} \quad (5)$$

Intuitively, the equilibrium tie-breaking probabilities ensure each firm has the same incentive to advertise the common upper utility bound \bar{u} . This creates significant flexibility and rules out any mass points in the distribution of advertised utilities. In a symmetric market, it follows that $x_a = x_b = 0.5$. More generally, x_i is determined in a way that offsets any asymmetries and aligns the incentives of the two firms to offer higher utilities at \bar{u} . In particular, as later verified in Section 4, x_i is i) decreasing in factors that discourage firm i from advertising higher utilities, such as firm i 's loyal share, θ_i , advertising costs, A_i , and monopoly profits, π_i^m , but ii) increasing in factors that encourage firm i to advertise higher utilities, such as firm i per-consumer profits at \bar{u} , $\pi_i(\bar{u})$. Now, as $x_a + x_b = 1$, one can sum (5) over $i = a, b$ and set equal to one:

$$\frac{A_a}{\pi_a(\bar{u})(1 - \theta_b) - \theta_a \pi_a^m} + \frac{A_b}{\pi_b(\bar{u})(1 - \theta_a) - \theta_b \pi_b^m} = 1 \quad (6)$$

This offers an implicit expression for the equilibrium upper utility bound, \bar{u} . The existence of a unique solution can be guaranteed if the advertising costs are sufficiently low. This follows as the LHS of (6) is strictly increasing in \bar{u} , A_a and A_b . Therefore, to ensure $\bar{u} > u^m$, an upper bound on advertising costs can be found by substituting u^m for \bar{u} in (6) and rearranging to give (7). This restriction is tighter than that provided under Assumption A.

$$\frac{A_a}{\pi_a^m} + \frac{A_b}{\pi_b^m} \leq 1 - \theta \quad (7)$$

The solution for \bar{u} further ensures a unique set of resulting interior tie-breaking probabilities x_a and x_b in (5) because it implies $\pi_i(\bar{u})(1 - \theta_j) + \theta_i \pi_i^m > 0$, and so each of the two elements on the LHS of (6) are bounded between zero and one.¹⁰

When advertising costs are too high to satisfy (7), the game has a different, simpler form of equilibrium. Here, for a relevant range of x_a and x_b , the high costs of advertising ensure that both firms simply select u^m and refrain from advertising. Proposition 1

¹⁰In an extreme case where the firms are asymmetric but $A_a = A_b \rightarrow 0$, the only way for the firms to share a common upper utility bound is for one firm to obtain all of the shoppers in the event of a tie. In particular, as detailed in Appendix B1, when $\pi_i^{-1} \left(\frac{\theta_i \pi_i^m}{1 - \theta_j} \right) < \pi_j^{-1} \left(\frac{\theta_j \pi_j^m}{1 - \theta_i} \right)$ the equilibrium has $x_i \rightarrow 0$, $x_j \rightarrow 1$, and $\bar{u} \rightarrow \pi_i^{-1} \left(\frac{\theta_i \pi_i^m}{1 - \theta_j} \right)$. Given this, firm j advertises with a probability converging to one, $\alpha_j \rightarrow 1$, but firm i still mixes over advertising with $\lim_{A \rightarrow 0} \alpha_i \in (0, 1)$. This limit equilibrium converges to the equilibrium of a model that allows for $A = 0$ explicitly without our tie-break rule. There, when $A = 0$, both firms find it optimal to advertise with probability one and to price on a common support. In equilibrium, the firms use equivalent utility distributions to those derived above, but firm i advertises u^m with a probability mass equivalent to $\lim_{A \rightarrow 0} (1 - \alpha_i)$.

formally summarizes our equilibrium results:

Proposition 1. *Under our tie-breaking rule, the game has the following unique equilibrium:*

1. If $\frac{A_a}{\pi_a^m} + \frac{A_b}{\pi_b^m} \leq 1 - \theta$, each firm i offers $u_i = u^m$ and does not advertise with probability $(1 - \alpha_i) \in (0, 1)$ according to (2), and advertises an offer u_i from the interval $(u^m, \bar{u}]$ according to (4) with probability α_i , where \bar{u} solves (6) and where $x_a = 1 - x_b$ is given by (5).
2. If $\frac{A_a}{\pi_a^m} + \frac{A_b}{\pi_b^m} \geq 1 - \theta$, both firms offer $u_i = u^m$ and never advertise, while shoppers visit firm a with a probability $x_a \in \left[1 - \frac{A_a}{\pi_a^m(1-\theta)}, \frac{A_b}{\pi_b^m(1-\theta)}\right]$.

Henceforth, we focus only on the more interesting equilibrium with sales behavior. In the next section, we provide a detailed discussion of the effects of market asymmetries by formally examining a selection of comparative statics. However, in the remainder of this section, we now briefly outline how our framework can reproduce and substantially extend a large range of equilibria from the existing literature through further specification of each firm's utility offer, u_i , and associated profits per consumer, $\pi_i(u_i)$. More precisely, while our assumptions sometimes differ to those used within the literature, we now demonstrate that our key equilibrium predictions for pricing, advertising, and purchase behavior are often identical to those predicted within the existing literature even in our more general setting.

Unit demand: Following our previous unit demand example, suppose $u_i = V_i - p_i$ and $\pi_i(u_i) = V_i - c_i - u_i$, where $u_i^m = 0$, and $\pi_i^m = V_i$. Under symmetry, this provides a simple clearinghouse equilibrium, and reproduces the (popularized) equilibrium of Varian (1980) when $A \rightarrow 0$. In particular, our results imply $x_i = 0.5$, such that we obtain the following familiar expressions: $\bar{\pi} = \frac{\theta(V-c)}{2} + A$, $1 - \alpha = \frac{2A}{(1-\theta)(V-c)}$, and $\bar{u} = \frac{2(1-\theta)(V-c)-4A}{2-\theta}$. By then using $F(p) = 1 - F(u)$, one can further derive $1 - F(p) = \frac{\theta(V-p)+4A}{2(1-\theta)(p-c)}$, with $\bar{p} = V - u^m = V$ and $\underline{p} = V - \bar{u} = c + \frac{\theta(V-c)+4A}{2-\theta}$.

Under asymmetry, the previous literature has largely focused on considering various combinations of asymmetries in terms of loyal shares, product values and/or costs under the restriction that $A_i = A_j \rightarrow 0$. Our framework synthesizes some versions of these equilibria by allowing for any θ_i , c_i , and V_i while also extending them to allow for positive

asymmetric advertising costs.¹¹ In particular, as further detailed in Appendix B1, the equilibrium with $A_i = A_j \rightarrow 0$ has $x_i \rightarrow 0$ and $x_j \rightarrow 1$ when $(1 - \theta_i)(V_i - c_i) < (1 - \theta_j)(V_j - c_j)$. By then denoting $\Delta V = V_i - V_j$, and noting that $F_i(u_j) = \Pr(u_i \leq u_j) = 1 - F_i(p_j + \Delta V)$ and $F_j(u_i) = 1 - F_j(p_i - \Delta V)$, we can derive some familiar expressions for $F_i(p)$ on $[V_i - \bar{u}, V_i]$ and $F_j(p)$ on $[V_j - \bar{u}, V_j]$, where firm j tends to always advertise, $\alpha_j \rightarrow 1$, but where firm i refrains from advertising with probability $1 - \alpha_i = 1 - F_i(V_i) \in (0, 1)$.

Downward-sloping demand: Suppose firm i has K_i products, where \mathbf{c}_i , \mathbf{p}_i and $\mathbf{q}_i(\mathbf{p}_i)$ denote the associated vectors of marginal costs, prices, and product demand functions for each consumer. The utility at firm i is then given by each consumer's associated surplus, $u_i = S(\mathbf{p}_i, \mathbf{q}_i(\mathbf{p}_i))$. To ensure Assumption U holds with $u_a^m = u_b^m$, we must restrict attention to cases with $\pi_a(u) = \pi_b(u) = \pi(u)$. Aside from settings where the firms have the same product range and costs $K_a = K_b$, $\mathbf{q}_a = \mathbf{q}_b$ and $\mathbf{c}_a = \mathbf{c}_b$, this can also be satisfied in some settings where demand and costs vary under specific relationships (as later detailed in Section 6). Given the sales equilibrium, each firm i then chooses its price vector to maximize its profits subject to supplying its required utility draw, u_i , with $\mathbf{p}_i^*(u_i) = \arg\max_{\mathbf{p}_i} \pi(\mathbf{p}_i)$ subject to $S(\mathbf{p}_i, \mathbf{q}(\mathbf{p}_i)) = u_i$. In a symmetric context, this set-up reproduces versions of i) clearinghouse equilibria (e.g. Baye et al. (2004a) and Baye et al. (2006)) or Baye and Morgan (2001) when firms sell single products given exogenous A and θ), and ii) the equilibrium of Simester (1997) when firms sell multiple products with zero marginal costs and $A \rightarrow 0$. More substantially, for any marginal costs, our framework extends these equilibria to permit positive asymmetric advertising costs, and asymmetric loyal-shares. See Appendix B2 for more formal details.

Other applications: Finally, our equilibrium can also be applied more broadly to explore some under-studied sales practices. For instance, through further specification of the utility and profit functions, current ongoing work by the authors examines i) two-part tariff sales as observed in markets such as energy and telecommunications, and ii) non-price sales where firms offer quantity extensions involving 'bonus packs' or 'X% Free'.

¹¹For example, among others, this application reproduces and extends i) the sales equilibrium of Narasimhan (1988) with vertically differentiated products and asymmetric loyal shares (pp.439-440), ii) the second stage equilibrium of Gu and Wenzel (2014)'s two-stage obfuscation game and the advertising game of Ireland (1993) and Roy (2000) which all allow for asymmetric θ_i , iii) Baye and De Vries (1992) trade equilibrium which allows for asymmetric c_i , iv) the second stage equilibrium of Schmidt (2013)'s two-stage cost-investment game which allows for asymmetric c_i and θ_i , and v) the second stage equilibrium of Jing (2007)'s two-stage quality-investment game which allows for asymmetric V_i , c_i and θ_i (Propositions 1 and 1A when loyal and shoppers share common preferences).

4 Comparative Statics

In this section, we use our framework to provide a range of comparative statics. For symmetric market cases, our findings extend the standard clearinghouse results to a more general setting. More substantially, for asymmetric market cases where the literature has previously offered a very limited understanding, we offer several new results that differ to the symmetric case in interesting and subtle ways.

4.1 Changes in Firms' Shares of Loyal Consumers

In the symmetric market case, one can verify a generalized form of the standard clearinghouse results - an increase in the proportion of loyal consumers, θ , deters firms from advertising better offers, and leads to a lower advertising probability, α , lower expected utility offers, $E(u)$, and higher equilibrium profits, $\bar{\pi}$. However, we now isolate the more complex effects from a change in an individual firm's loyal share, θ_i . As these effects are difficult to characterize generally, we focus on evaluating the asymmetric comparative statics at the point of symmetry. To proceed, one must also stipulate whether the increase in θ_i comes at the expense of a reduced proportion of rival loyal, $\theta_j = \theta - \theta_i$, or a reduced proportion of shoppers, $1 - (\theta_i + \theta_j)$. We first consider the former:

Proposition 2. *Consider an increase in firm i 's loyal share θ_i (and associated reduction in firm j 's loyal share) in an otherwise symmetric market. Starting from $\theta_i = \theta_j$, this i) decreases x_i , ii) increases $\bar{\pi}_i$, iii) decreases $\bar{\pi}_j$, iv) decreases α_i and $E(u_i)$, and v) increases α_j and $E(u_j)$.*

Following an increase in θ_i , firm i is less willing to advertise higher utilities holding all else constant. Therefore, to maintain a common \bar{u} , this must be offset by a reduction in firm i 's tie-break share, x_i . Hence, an increase in θ_i produces both a direct effect, and an indirect effect through the reduction in x_i . First, take firm i 's profits $\bar{\pi}_i = \theta_i \pi^m + \left(\frac{x_i}{1-x_i}\right)A$. The positive direct effect dominates the negative indirect effect to ensure that $\bar{\pi}_i$ rises. In contrast, firm j 's profits fall despite the positive indirect effect in raising x_j because it also receives a reduction in θ_j . Second, the indirect effect's impact in reducing x_i and increasing x_j ensure that firm i 's advertising probability decreases and that firm j 's increases. Finally, in terms of the utility distributions, these changes prompt lower average utility offers at firm i , but higher average offers at firm j . While very intuitive, this last result differs to the equivalent finding in Arnold et al. (2011) who allow for asymmetric θ_i with unit demand and $A > 0$ under an equal tie-break rule, $x_i = 0.5$. Instead, they

suggest that the firm with the larger loyal share, firm i , adopts a pricing strategy when advertising that is *more* aggressive than its rival. This contrasts sharply to the results in our model, and in models with $A = 0$ such as Narasimhan (1988).¹²

4.2 Changes in Firms' Advertising Costs

As before, one can verify a generalized form of the standard clearinghouse results in the symmetric market case - an increase in advertising costs, A , deters the firms from advertising better offers and softens competition for the shoppers in a way which lowers the firms' advertising probability, α , raises equilibrium profits, $\bar{\pi}$, and reduces consumers' expected utility offers, $E(u)$. We now document some new, more nuanced, effects from a change in an individual firm's advertising cost, A_i .

Proposition 3. *Consider an increase in A_i in an otherwise symmetric market. This leads to i) a lower x_i , ii) no change in $\bar{\pi}_i$, iii) a one-for-one increase in $\bar{\pi}_j$, and iv) a reduction in advertising probabilities and expected utility offers for both firms.*

Ceteris paribus, an increase in A_i reduces the incentives for firm i to advertise higher utilities. Therefore, this must be offset by a reduction in firm i 's tie-break share, x_i , to maintain a common \bar{u} . In particular, in this setting, it follows that $x_i = \frac{A_j}{A_i + A_j}$ such that the firm with the higher advertising cost receives the smaller share of shoppers in a tie. Firm i 's profits can then be written as $\bar{\pi}_i = \frac{\theta}{2}\pi^m + (\frac{A_j}{A_i})A_i$ where an increase in own advertising costs has no aggregate effect because the direct positive effect is exactly offset by the negative indirect effect through the reduction in x_i . However, an increase in A_i increases the rival firm j 's profits via the indirect effect in raising x_j . Hence, after decomposing the effects in this way, we can see that it is rival rather than own advertising costs that matter for individual profits. Finally, given the otherwise symmetric market, Proposition 3 also suggests that both firms employ an identical advertising probability and an identical utility distribution which depend only on total advertising costs, $A_a + A_b$. Thus, any increase in A_i generates a common reduction in advertising probabilities and expected utility offers for both firms.¹³

¹²With two exceptions, our findings remain robust in the alternative case where the increase in firm i 's share of loyals comes at the expense of a reduction in the proportion of shoppers. First, an increase in θ_i now raises firm j 's profits because it only benefits from the rise in x_j without the reduction in θ_j . Second, an increase in θ_i can provide reversed effects on α_j and $E(u_j)$ if advertising costs are relatively high. This highlights the importance of considering positive advertising costs and arises because of the conflicting effects from a decrease in the proportion of shoppers, $(1 - \theta)$, and an increase in x_j which varies in the level of A . However, it remains that firm i still offers a lower average utility than firm j . (Full details on request).

¹³These last results become less stark in our later analysis when $n > 2$ where advertising cost asymme-

4.3 Changes in Firms' Profit Functions

Firms' profit functions may differ due to many factors including costs, products, demand, or pricing technologies. As discussed in Section 3, the literature has only been able to consider a few such asymmetries under the case of zero advertising costs, and has certainly been unable to provide any general set of comparative statics. Some related technical difficulties are also present in our framework because such asymmetries may i) influence profits differently at different levels of utility, and ii) lead to asymmetric levels of monopoly utility, $u_i^m \neq u_j^m$.

However, we can consider how the effects of asymmetries in firms' profit functions for a tractable case where the shape of the profit function is always preserved and u_i^m remains constant. This 'multiplicative' case involves $\pi_i(u) = (1 - \tau_i)\pi(u)$ with $\tau_i \in [0, 1)$, where an increase in τ_i represents a reduction in firm i 's per-consumer profits or 'profitability' via any associated change in costs, products, demand or pricing technology.

Proposition 4. *Suppose $\pi_i(u) = (1 - \tau_i)\pi(u)$ and consider an increase in τ_i within an otherwise symmetric market. This leads to i) a lower x_i , ii) a lower $\bar{\pi}_i$, iii) a higher $\bar{\pi}_j$, and iv) a reduction in advertising probabilities and expected utility offers for both firms.*

Surprisingly, these results share some features with those for asymmetric advertising costs. This arises because many of the equilibrium expressions depend only on the ratio $(1 - \tau_i)/A_i$. For instance, similar to the previous subsection, $x_i = \frac{1-\tau_i}{2-\tau_i-\tau_j}$. Intuitively, an increase in τ_i reduces firm i 's incentives to advertise higher utilities and must therefore be offset by a decrease in x_i . Also like the case of asymmetric advertising costs, any increase in τ_i generates an equal reduction in advertising probabilities and expected utility offers for both firms. However, the two cases differ when considering equilibrium profits, $\bar{\pi}_i = \frac{\theta}{2}(1 - \tau_i)\pi^m + \frac{1-\tau_i}{1-\tau_j}A$. An increase in τ_i , reduces $\bar{\pi}_i$ both directly, and indirectly via an decrease in x_i , while an increase in τ_j , increases $\bar{\pi}_i$ via an indirect effect in increasing x_i .

5 More Than Two Firms

In this section, we now illustrate the framework's further ability to offer new results by analyzing the number and type of firms that engage in sales behavior in markets with more than firms, $n > 2$.

tries begin to influence advertising probabilities and utility distributions. However, the direction of the reported statics remain.

The sales literature with heterogeneous firms is particularly scant when $n > 2$ because existing models quickly become complex and intractable. However, in a seminal paper, Baye et al. (1992) establish a central result that states that only two firms can ever engage in sales behavior when firms vary in their loyal shares within a unit-demand clearinghouse model with zero advertising costs. Intuitively, the remaining firms with relatively large loyal shares are less willing to lower their price and prefer, instead, to simply price highly to their loyal consumers. This finding has been extended to allow the firms to vary in their product values (Kocas and Kiyak (2006)) or costs (Shelegia (2012)). It has also been used as the foundation for a number of studies, such as those aiming to endogenize consumer loyalty (e.g. Chioveanu (2008)).

In contrast, we now show that this two-firm result is a special case of our more general findings. In particular, once we allow for positive advertising costs, any number of firms $k^* \in [2, n]$ can engage in equilibrium sales and advertising behavior, and, paradoxically, this number is (weakly) increasing in the level of advertising costs. Further, we can then use the flexibility of our framework to provide a broad characterization of the types of firms that are likely to advertise in equilibrium.

Compared to our previous duopoly model, there are two potential sources of equilibrium multiplicity. First, similar to an insight by Baye et al. (1992) for zero advertising costs, the equilibrium distributions and supports may no longer be unique even in a symmetric setting. In particular, provided at least two firms mix in any given interval within $(u^m, \bar{u}]$, there may be equilibria where other firms do not mix within that region. Second, there is a new form of multiplicity that is specific to our framework. Indeed, it is no longer true that each firm's tie-break share, x_i , is uniquely defined in equilibria where one or more firms never advertise, $\alpha_i = 0$ for some i .

Therefore, in addition to our previous assumptions, we choose to focus on equilibria under the following two restrictions. First, with no loss to our qualitative predictions, we focus on equilibria where within any tie, shoppers disregard firms that *never* advertise, such that $x_i = 0$ if $\alpha_i = 0$. Second, for tractability, we focus on equilibria where any advertising firm advertises over the same convex support $(u^m, \bar{u}]$, such that $F'_i(u) = f_i(u) > 0$ for all $u \in (u^m, \bar{u}]$ if $\alpha_i > 0$.¹⁴

In what follows, we denote $\theta_{-i} = \theta - \theta_i$ as the total share of loyal consumers that are

¹⁴These restrictions may be less necessary within a symmetric setting. Indeed, within a symmetric n -firm clearinghouse model with unit demand and positive advertising costs, Arnold and Zhang (2014) show that the symmetric equilibrium is unique and that asymmetric equilibria do not exist. One may be able to use similar methods in the symmetric setting of our more general framework. However, to be sure, we apply our restrictions to both the symmetric and asymmetric settings.

not loyal to firm i , and k^* as the number of firms that advertise in equilibrium. We also denote \tilde{u}_i as the highest utility that firm i could possibly be willing to advertise, where:

$$\tilde{u}_i \equiv \pi_i^{-1} \left(\frac{\theta_i \pi_i^m + A_i}{1 - \theta_{-i}} \right) \quad (8)$$

This derives from equating firm i 's highest possible profits from advertising $\pi_i(\tilde{u}_i)(1 - \theta_{-i}) - A_i$, with its lowest possible profits from not advertising, $\theta_i \pi_i^m$. Without loss, we then index the firms in (weakly) decreasing order of \tilde{u}_i from 1 to n and focus on two settings: i) a quasi-symmetric setting where $\tilde{u}_i = \tilde{u} > u^m$ for all i , and ii) a strict asymmetric setting where $\tilde{u}_1 > \tilde{u}_2 > \dots > \tilde{u}_n > u^m$ such that firm n is the least willing to advertise high utility levels.¹⁵ We first state a useful preliminary result.

Lemma 2. *In any equilibrium that satisfies our restrictions with common upper utility bound \bar{u} , firm i advertises if and only if $\tilde{u}_i \geq \bar{u}$. Hence, i) if $k^* = n$ then $\tilde{u}_n \geq \bar{u}$, and ii) if $k^* \in [2, n)$ then $\tilde{u}_{k^*} \geq \bar{u} > \tilde{u}_{k^*+1}$.*

This follows by contradiction. Recall that our equilibrium restrictions require firms that never advertise to receive zero shoppers in a tie, $x_i = 0$ if $\alpha_i = 0$. By definition of \tilde{u}_i , any firm i with $x_i = 0$ makes the same profit from not advertising as it would if it advertised \tilde{u}_i and won all the shoppers. Therefore, any firm with $\tilde{u}_i > \bar{u}$ that never advertises with $\alpha_i = 0$ would always prefer to advertise \bar{u} . Similarly, any firm with $\tilde{u}_i < \bar{u}$ that does advertise with $\alpha_i > 0$ would always prefer to refrain from advertising given our restriction that any advertising firm must advertise over the entire equilibrium support $u \in (u^m, \bar{u}]$.

Using Lemma 2, we now derive the game equilibria under our restrictions. While we show that the equilibrium will always be unique, it is hard to demonstrate existence for the general case without specifying exact profit functions. However, existence can be shown for special cases including when the firms are (sufficiently) symmetric, or differ only in their advertising costs.

Proposition 5. *When an equilibrium exists under our restrictions, it is unique. In such an equilibrium, firms $i = \{1, \dots, k^*\}$ advertise with interior probability over $(u^m, \bar{u}]$, while any remaining firms, $i = \{k^* + 1, \dots, n\}$ set $u_i = u^m$ and never advertise, where k^* is*

¹⁵A third setting where some firms have the same \tilde{u} but some firms do not can also be analyzed. However, this brings unnecessary complications and so we omit it for brevity.

uniquely defined by

$$k^* = \begin{cases} n & \text{if } \sum_{i=1}^n \frac{A_i}{h_i(\tilde{u}_n)} > (n-1) \geq \sum_{i=1}^n \frac{A_i}{(1-\theta)\pi_i^m} \\ k \in [2, n) & \text{if } \sum_{i=1}^k \frac{A_i}{h_i(\tilde{u}_k)} > (k-1) \geq \sum_{i=1}^{k+1} \frac{A_i}{h_i(\tilde{u}_{k+1})} - 1 \end{cases} \quad (9)$$

and where $h_i(u) \equiv \pi_i(u)(1 - \theta_{-i}) - \theta_i \pi_i^m > 0$.

Note that $\frac{A_i}{h_i(u)} \leq 1$ for $u \in [u^m, \tilde{u}_i]$ and $\frac{A_i}{h_i(u_i)} = 1$, so the last term in both of the sums on the second line of (9) equals 1. This means that in the second line if the second inequality holds for $k = 2$, the first inequality is automatically satisfied, as consistent with our results for $n = 2$.

To begin to understand Proposition 5, first consider the quasi-symmetric setting with $\tilde{u}_i = \tilde{u} \geq u^m$ for all i . By using (9), the only possible equilibrium under our restrictions then has $k^* = n$. As detailed in the proof, this equilibrium is symmetric and resembles that under duopoly: all firms engage in sales behavior on the common interval $[u^m, \bar{u}]$. In general, even in this case, neither the utility distributions, $F_i(u)$, or advertising probabilities, α_i , will be the same for all the firms unless the firms are exactly symmetric. As under duopoly, this equilibrium requires advertising costs to be sufficiently small, $(n-1) \geq \sum_{i=1}^n \frac{A_i}{(1-\theta)\pi_i^m}$. However, the other corresponding condition on advertising costs in (9) does not bind (as further explained below).

For the rest of this section, we consider the asymmetric setting where $\tilde{u}_1 > \tilde{u}_2 > \dots > \tilde{u}_n > u^m$. Here, Proposition 5 provides a number of results to help understand both the number and the type of firms that use sales and advertising in equilibrium. We consider each in turn.

5.1 The Number of Firms that Engage in Sales

The number of firms that use sales in equilibrium, k^* , is uniquely determined by the conditions in (9). First, consider the expressions on the right-hand-side. These provide an upper bound on advertising costs. When $k^* = n$, the upper bound ensures that \bar{u} is sufficiently large, with $\bar{u} \geq u^m$, such that all firms have an incentive to advertise. In contrast, when $k^* < n$, the upper bound has a different role. It ensures that \bar{u} is sufficiently large, with $\bar{u} > \tilde{u}_{k^*+1}$, such that all non-advertising firms have no incentive to advertise. Now consider the expressions on the left-hand-side of (9). These provide a lower bound to advertising costs for any $k^* \in [2, n]$. The lower bound ensures that \bar{u} is sufficiently small, with $\bar{u} \leq \tilde{u}_{k^*}$, such that all advertising firms are willing to advertise

$u_i = \bar{u}$ without requiring $x_i < 0$. When $n = 2$, or in the quasi-symmetric case when $n > 2$, this lower bound does not bind because all x_i are guaranteed to be positive in equilibrium. However, in the asymmetric setting with $n > 2$, this lower bound becomes important to ensure each x_i is well-defined and positive.

To provide a more intuitive explanation of k^* , first suppose the firms vary in \tilde{u} but share a common advertising cost, $A_i = A$ for all i . Notice that \tilde{u}_i depends on A , thus we assume here that the ranking of \tilde{u}_i does not change with A . Using Proposition 5, we then note the following result.

Corollary 1. *Suppose $\tilde{u}_1 > \tilde{u}_2 > \dots > \tilde{u}_n \geq u^m$ but $A_i = A$ for all i . In the limit, $k^* = 2$ when $A \rightarrow 0$. In addition, if the firms have sufficiently symmetric profit functions, with $\sum_{i=1}^n \frac{1}{\pi_i^m} > \frac{n-1}{\pi_n^m}$, then $k^* = n$ when $A \rightarrow \frac{(n-1)(1-\theta)}{\sum_i 1/\pi_i^m}$.*

When $A \rightarrow 0$ our findings confirm and reproduce the result in the existing literature that only two firms can engage in sales behavior (Baye et al. (1992), Kocas and Kiyak (2006), Shelegia (2012)). However, while these papers allow for forms of heterogeneity under unit demand, our results extend the two-firm result to a far broader range of market settings. Intuitively, when A goes to zero, the only way to have any firms advertise up to a common utility upper bound is to give the firm with the highest incentive to advertise, firm 1, all the shoppers in a tie in order to discourage it from advertising higher utilities. As this can only be done for one firm to ensure that firm 1 and firm 2 are on an equal footing, no further firms can engage in sales behavior.

However, for larger levels of A , our findings suggest that this two-firm result is a special case of a new and more general relationship. Indeed, once we allow for positive advertising costs, any number of firms $k^* \in [2, n]$ may engage in equilibrium sales and advertising behavior. In order to allow all n firms to advertise, A has to be sufficiently high, and firms' profit functions have to be sufficiently similar. For the latter, the condition in the proposition requires that the sum of the inverse of monopoly profits for all firms is higher than the inverse for the firm with lowest monopoly profits (firm n) multiplied by $n - 1$.¹⁶ This condition is trivially satisfied for 2 firms, and becomes increasingly stringent as n grows. As $n \rightarrow \infty$ it requires that firms are arbitrarily symmetric in their monopoly profits. However, the condition makes no restrictions on the firms' loyal shares, thus any disparity there can be alleviated by sufficiently high advertising costs and appropriately

¹⁶Holding π_n^m fixed, this condition is most stringent when $\pi_i^m = \pi_1^m$ for all $i \leq n-1$. Thus the sufficient condition is $\pi_n^m > \frac{n-2}{n-1}\pi_1^m$. This is always satisfied for symmetric firms or when $n = 2$. For $n = 3$, the sufficient condition requires that the most profitable firm, firm 1, is less than twice as profitable as the least profitable firm, firm n .

set tie-break shares, x_i .

In most cases where the condition on monopoly profits is satisfied, this leads to a surprising result - the equilibrium number of advertising firms, k^* , will weakly increase as advertising costs, A , increases from 0 to $\frac{(n-1)(1-\theta)}{\sum_i 1/\pi_i^m}$.¹⁷ Provided that the firms have sufficiently symmetric monopoly profits, this relationship arises because, \bar{u} (as implicitly defined by (14)) converges to u^m faster than \tilde{u}_n , allowing for the possibility that all firms may advertise up to the common utility upper bound. More intuitively, an increase in A softens advertising competition for shoppers and shrinks the differences between \tilde{u}_i amongst firms (as all converge to u^m). This allows for a redistribution of the tie-break shares, x_i , in such a way that all firms share a common upper utility bound. This is not true for lower levels of A because the spread between firms' high utility advertising abilities is too large, and there are not enough tie-break shares to balance each firm's incentives. In the extreme where $A \rightarrow 0$, only firm 2 can be made competitive with firm 1 (by allocating all the shoppers to it in case of a tie), and all remaining firms refrain from sales competition.

5.2 The Types of Firms that Engage in Sales

Finally, having established k^* , we now examine which types of firms are likely to engage in sales and advertising behavior. Existing results within the literature only consider some specific dimensions under unit demand and zero advertising (e.g. Baye et al. (1992), Kocas and Kiyak (2006), Shelegia (2012)). However, in our general setting, we can offer a broader characterization. In particular, when $k^* < n$, Proposition 5 implies that the firms using sales and advertising will be the firms with the lowest values of \tilde{u}_i , firms $i = \{1, \dots, k^*\}$. Due to the implicit definition of \tilde{u}_i , we restrict our attention to the multiplicative profit function that was considered in Section 4.3. Using (8), we can then state:

Corollary 2. *Suppose $k^* < n$. Ceteris paribus, the firms that engage in sales and advertising behavior will be the firms with the lowest A_i and θ_i , and the highest $\pi_i(u)$.*

Intuitively, firms with the lowest advertising costs and loyal shares have the highest incentives to advertise high utilities, and will therefore be most likely to engage in sales in equilibrium. When profits are multiplicative, \tilde{u}_i is increasing in the profit level. Thus, more profitable firms are also most likely to engage in sales.

¹⁷Notice that \tilde{u}_i depend on A , as well as does \bar{u} , so k^* may be non-monotonic as A transitions from 0 to $\frac{(n-1)(1-\theta)}{\sum_i 1/\pi_i^m}$, but the overall increase is guaranteed.

6 Implications for Empirical Research

Understanding the effects of firm heterogeneity in markets is a necessary challenge for empirical work. However, without the foundation of an adequate theoretical model, typical empirical papers on sales behavior or price dispersion are forced to resort to a ‘cleaning’ procedure that is known to be restrictive. In this section, we now illustrate how our framework can be used to better understand when such an approach is theoretically valid, and to provide the basis for a modified methodology for settings where the cleaning procedure is not valid.

Empirical studies often find that firms employ sales that are consistent with the use of mixed strategies, but exhibit differences in their average price levels.¹⁸ This pattern is driven by two forms of price dispersion. The first ‘spatial’ form of price dispersion arises from inter-firm differences that remain over time, such as those arising from differences in firms’ characteristics, products, or costs. The second ‘temporal’ form involves price differences that vary over time, including those generated by sales behavior. Empirical studies often wish to separate these two forms of price dispersion. In particular, to focus only on the temporal form, empirical papers typically ‘clean’ their raw price data by retrieving a set of price residuals from a price regression involving observable firm characteristics or firm-level fixed effects. The price residuals are then interpreted as resulting from a homogeneous symmetric market and used to either i) perform descriptive/reduced-form analysis of the features of temporal price dispersion, or ii) conduct structural estimations of market parameters.¹⁹

However, this procedure is known to be restrictive. First, Lewis (2008) notes that both the firm characteristics and their impact on the firms’ profit functions must be stable over time. Second, Chandra and Tappata (2011) suggest that the interaction between the firm-level fixed effects and the source of (price) dispersion must be additively separable. These arguments are extended by Wildenbeest (2011) who provides the only formal justification for the procedure within the literature.

A version of Wildenbeest’s arguments can be derived within our n -firm clearinghouse framework, where in contrast, we generalize to positive advertising costs. Suppose consumers have unit demand. Firms can vary in quality and costs subject to a common

¹⁸e.g. Lach (2002), Baye et al. (2004a), Baye et al. (2004b), Kocas and Kiyak (2006), Lewis (2008), Chandra and Tappata (2011), Wildenbeest (2011), Giuliatti et al. (2014) and Pennerstorfer et al. (2015).

¹⁹Examples of group i) include Sorensen (2000), Lach (2002), Brown and Goolsbee (2002), Barron et al. (2004), Lewis (2008), Chandra and Tappata (2011), Pennerstorfer et al. (2015), and Sherman and Weiss (2015). Examples of group ii) include Wildenbeest (2011), Moraga-González et al. (2013), Giuliatti et al. (2014), Allen et al. (2014), and An et al. (2015).

value-cost markup, $V_i - c_i = \Psi$ for all i such that the firms share a symmetric profit function, $\pi_i(u) = V_i - c_i - u = \Psi - u$. Assuming symmetric loyal shares and advertising costs, the firms then employ a symmetric equilibrium utility distribution with $u^m = 0$, $\bar{u} = \frac{n\Psi(1-\theta)-2nA}{n-(n-1)\theta}$ and $F(u) = \frac{\theta u + 2nA}{n(1-\theta)(\Psi-u)}$. Under unit demand, $p_i(u_i) = V_i - u_i$. Therefore, give the symmetric utility distribution, the firm's subsequent price distributions are simple iid translations of each other. Under the assumption that the firms play a stationary repeated game of finite horizon, each firm then chooses its utility level for each period as a draw from the equilibrium distribution. Now suppose that the econometrician observes a panel of price observations for each firm. The econometrician can then obtain measures of the firms' utility offers that are entirely cleaned of the effects of firm-heterogeneity by using one of two possible methods. First, one can use the observed maximum price for each firm to infer V_i directly, and simply adjust the observed prices to recover $u_i = V_i - p_i$. Second, given $p_i = V_i - u_i$, one can regress the raw price data on a set of firm-level fixed effects, $p_{it} = \alpha + \delta_i + \varepsilon_{it}$, to return a set of 'cleaned' price residuals that correctly proxy the utility draws up to a positive constant.²⁰ As the first method may be subject to data outliers, the literature typically employs the second method.

Wildenbeest's justification for the cleaning procedure applies only for settings where the firms employ a symmetric utility distribution and where consumers have unit demand. By using our more general theoretical framework, we now consider some alternative settings to further understand the validity of the cleaning procedure and to offer some modified procedures.

6.1 Downward-Sloping Demand

To ensure that the firms employ a symmetric utility distribution in a market with symmetric loyal shares and advertising costs, our framework suggests that firms need to share a common profit function, $\pi(u)$. Under unit demand, this was guaranteed by Wildenbeest's constant value-cost markup assumption. However, if instead, demand is downward-sloping, we require a different condition. To illustrate, suppose each firm has a linear per-consumer demand function that varies only in its intercept, $q_i(p) = a_i - bp$ where $a_i \geq 0$ and $b > 0$. Further suppose that firm i has marginal cost $c_i \geq 0$. One can then use the results from Section 3 and Appendix B2 to show that $u = \frac{(a_i - bp)^2}{2b}$ and

²⁰In more detail, the estimated residuals $\hat{\varepsilon}_{it} \equiv p_{it} - p_i^{ave}$ where p_i^{ave} is the average price chosen by firm i . Given unit demand and the symmetric equilibrium utility distribution with average utility offer, u^{ave} , it follows that $p_{it} = V_i - u_{it}$ and $p_i^{ave} = V_i - u^{ave}$, such that the estimated residuals provide a negative measure of the utility draws, $\hat{\varepsilon}_{it} \equiv -(u_{it} - u^{ave})$.

$\pi_i(u) = \frac{1}{b}[a_i - bc_i - \sqrt{2bu}][\sqrt{2bu}]$ such that the firms have a symmetric profit function if and only if $a_i - bc_i = \Psi$ for all i . Intuitively, this more general condition captures some sense of Wildenbeest’s constant value-cost assumption.

Given this condition, one would then hope to be able to recover the firms’ utility draws from the raw price data. However, unlike unit demand, the relationship between prices and utilities is non-linear, $u = \frac{(a_i - bp)^2}{2b}$. Therefore, despite the possibility of a symmetric utility distribution, the literature’s cleaning procedure cannot be applied under downward-sloping demand. Instead, one would have to implement a more complex and data-intensive procedure to recover the utility draws by using additional quantity data to estimate each firm’s demand function.

6.2 Asymmetric Utility Distributions

To further explore the cleaning procedure, we now return to the case of unit demand, but depart from Wildenbeest’s set-up by considering a setting where the firms use asymmetric utility distributions, $F_i(u) \neq F(u)$. Here, the fixed-effects procedure is no longer valid because the firms’ price distributions are not simple iid translations of each other.

However, by using the insights of our model, one could consider the following modified procedure. Given $F_i(u) \neq F(u)$, one has to recover the utility draws separately for each firm. Instead of using the fixed-effects regression, one may think about doing this by estimating a set of firm-specific price regressions. However, this is not valid either because the interpreted residuals from each regression, $\hat{\epsilon}_{it} \equiv p_{it} - p_i^{ave} \equiv u_{it} - u_i^{ave}$, cannot be compared as the firms now differ in their average utility levels, u_i^{ave} , due to differences in their loyal shares, or value-cost margins. Instead, one would have to recover the utility draws separately under the more direct method. In particular, by using the observed maximum price of each firm, one would infer V_i , and then calculate each firm’s utility offer using $u_i = V_i - p_i$. As this approach may be sensitive to data outliers, a more robust method may be to infer V_i by establishing each firm’s ‘regular’ price instead, using a procedure such as that used in Hosken and Reiffen (2004).

Having recovered the utility draws, one could then use our theoretical insights to analyse the observed price dispersion, or to estimate a structural model within the asymmetric market setting. For instance, by using our theoretical predictions for the utility distribution and perhaps the advertising probability, one may be able to use price and advertising data to estimate each firm’s loyal share, θ_i , and advertising cost, A_i .

7 Conclusions

Due to the potential technical complexities, the existing clearinghouse sales literature has been unable to fully consider the effects of firm heterogeneity on sales competition. This has restricted theoretical and empirical understanding, and limited policy guidance in these important markets. Indeed, as Baye and Morgan (2009) state (p.1151), “...*little is known about asymmetric models within this class. Breakthroughs on this front would not only constitute a major theoretical advance, but permit a tighter fit between the underlying theory and empirics*”.

The current paper has tried to fill this gap by providing a fully asymmetric version of the clearinghouse sales framework. The framework can i) reproduce and extend many equilibria from the existing literature, ii) offer a range of new results on how firm heterogeneity affects sales behavior and market performance, iii) provide original insights into the number and type of firms that engage in sales behavior, and iv) offer a deeper foundation to assess and extend current empirical procedures.

Our framework should induce future research in at least two wider respects. First, it should encourage a new range of empirical work. Aside from providing a possible new procedure for analysing sales data under firm heterogeneity, our results also make some new predictions that are yet to be tested, such as the prediction that the number of firms using sales can increase under higher advertising costs. Second, future work could extend our framework to analyse a wider set of sales practices beyond those considered here under our clearinghouse assumptions. For instance, a modified version of our framework could consider sales within dynamic consumer inventory models (e.g. Hong et al. (2002)), or explicit consumer search models (e.g. Stahl (1989) or Janssen and Moraga-González (2004)).

Appendix A - Main Proofs

Proof of Lemma 1. The fact that there are no gaps follows from the following: if there was a gap, it would have to have been common between both firms. Either firm could then move the mass of advertised utilities from just above the gap to the bottom of the gap. This would not alter the probability that those utilities were the best, but would make the firm earn higher profits. Similarly, if there were point masses in equilibrium, there would either be ties with a positive probability, in which case each firm could increase its profits by avoiding them, or only one of the firms would have a point mass, in which case

the other firm could increase its profits by moving the mass of its utility offers from just below the point mass to just above. Finally, the lower bound of advertised utilities has to be u^m , or otherwise firms could move the mass of lowest advertised utilities down to u^m without reducing the probability they capture shoppers and increasing profits. \square

Proof of Proposition 1. Part 1. When A_i -s satisfy the condition, we know that \bar{u} exists that solves (6). By construction, there is an equilibrium where firms advertise with probability given in (2) and set $u = u^m$, advertise utilities in $(u^m, \bar{u}]$ according to F in (4). Profits are equal for all utility offers and advertising strategies in equilibrium. Not advertising and setting $u \neq u^m$ can never be profitable because when not advertising demand is always the same, but per consumer profit is maximized at u^m . Advertising u outside of $(u^m, \bar{u}]$ gives strictly lower profit than advertising inside the interval. For $u < u^m$ that is because no shoppers are attracted, and per consumer profit is lower than at u^m . For $u > \bar{u}$, all shoppers are attracted just as with \bar{u} , but per consumer profit is lower. Thus equilibrium is described in Part 1 is correct.

What remains to be shown is that no other equilibrium exists for these values of A_i . There are three other possibilities. First, firms do not advertise at all; second, only one firm ever advertises; third, both firms advertise but not in the way described in the proposition above.

Consider the first possibility. There has to be some allocation of loyals to firm i in equilibrium (x_i). Regardless of this allocation, both firms will set $u = u^m$. For this to be an equilibrium, no firm has to have profitable deviation. Each firm can deviate to advertising utility slightly above u^m and capture all the loyals. Firm i gains $(1 - \theta)(1 - x_i)\pi_i^m$ by doing so, therefore in equilibrium we should have $(1 - \theta)(1 - x_i)\pi_i^m \leq A_i$ for $i = a, b$. This can be rewritten as $x_i \geq 1 - \frac{A_i}{(1 - \theta)\pi_i^m}$. Summing up over i gives $1 \geq 2 - \sum_i \frac{A_i}{(1 - \theta)\pi_i^m} \implies \sum_i \frac{A_i}{\pi_i^m} \geq (1 - \theta)$, a contradiction. Therefore, if A_i satisfy the condition, it is not possible that neither firm advertises.

Now consider the second possibility. Assume firm i advertises and firm j does not. Regardless of what firm i does, firm j 's optimal non-advertising strategy is to set $u = u^m$, so in any such equilibrium it has to do so. Given this, and that firm j advertises with positive probability, firm b should advertise utility slightly above u^m and capture all the shoppers (this in itself is not well defined, but we shall find contradiction nevertheless). Firm i can only be prevented from also advertising even slightly higher utility if $(1 - \theta_j)\pi_i^m < A_i$, contradiction with Assumption A.

Finally, the third possibility is that both firms advertise with positive probability, but not according to our equilibrium. If so, following standard arguments, they have to advertise with mixed strategies in a common utility interval without point masses. The lower bound of this interval, \underline{u} , has to be u^m . Assume the opposite. Clearly $\underline{u} < u^m$ cannot be in equilibrium, as $u = u^m$ strictly dominates (attracts shopper with higher probability, and increases per consumer profit). If $\underline{u} > u^m$, then given that there can be no point masses, advertising any u in the interval $[u^m, \underline{u})$ strictly dominates \underline{u} because shoppers are attracted with the same probability, but per consumer profit is higher. Thus, if firms are to advertise, they will do so on the interval $(u^m, \bar{u}]$. The left bound is open because advertising u^m is always dominated by not advertising and setting $u = u^m$. The only way such an equilibrium differs from our equilibrium, is that $\alpha_a = \alpha_b = 0$. This is impossible for any $A_i > 0$ ($i = a, b$), because profit from advertising slightly above u^m would be given by $\pi_i^m(1 - \theta_j) - A_i$, and has to be no smaller than $\pi_i^m\theta_i$, profit from deviating to not advertising. This implies $A_i \geq (1 - \theta)\pi_i^m$, which contradicts $\frac{A_a}{\pi_a^m} + \frac{A_b}{\pi_b^m} \leq 1 - \theta$ and Assumption A. This completes proofs of Part 1.

Part 2 can be proven as follows. If neither firm advertises, and shoppers are allocated according to x_a in the interval $x_a \in \left[1 - \frac{A_a}{\pi_a^m(1-\theta)}, \frac{A_b}{\pi_b^m(1-\theta)}\right]$, no firm would want to deviate and advertise, because even advertising slightly above u^m would not be profitable (x_a is chosen in such a way). Thus this is an equilibrium, so we need to show that no other equilibrium exists. Two possibilities exist, one where only one firm advertises, and one where both firms advertise. The former cannot be the case by Assumption A (see proof of Part 1 above). The latter is not possible because if both firms advertise, they have to do so in the fashion described in Part 1 of the proposition, and such equilibrium cannot be constructed because A exceeds the threshold. \square

Proof of Proposition 2. Let $\pi_i(u) = \pi(u)$, $A_i = A$ and $\theta_j = \theta - \theta_i$. From (6), $\frac{\partial \bar{u}}{\partial \theta_i} = 0$ after we impose symmetry ex post with $\theta_i = \theta_j$. By using this with the derivative of (5), we gain $\frac{\partial x_i}{\partial \theta_i} = -\frac{A[\pi^m - \pi(\bar{u})]}{[\pi(\bar{u})(1 - (\theta/2)) - (\theta/2)\pi^m]^2} < 0$. These two results can then be used to help find the relevant derivatives: i) Follows directly from above. For ii) and iii), using (3) gives $\frac{\partial \bar{\pi}_i}{\partial \theta_i} = \pi(\bar{u}) > 0$ and $\frac{\partial \bar{\pi}_j}{\partial \theta_i} = -\pi(\bar{u}) < 0$. For iv) and v), using (2), $\frac{\partial \alpha_i}{\partial \theta_i} = -\frac{[\pi^m - \pi(\bar{u})]}{(1 - \theta)\pi^m} < 0$, and $\frac{\partial \alpha_j}{\partial \theta_i} = \frac{\pi^m - \pi(\bar{u})}{(1 - \theta)\pi^m} > 0$. Further, from (4), $\frac{\partial F_i}{\partial \theta_i} = \frac{\pi(u) - \pi(\bar{u})}{(1 - \theta)\pi(u)} > 0$ and $\frac{\partial F_j}{\partial \theta_i} = -\frac{\pi(u) - \pi(\bar{u})}{(1 - \theta)\pi(u)} < 0$ such that expected utility at firm i (firm j) decreases (increases). \square

Proof of Proposition 3. Given $\pi_i(u) = \pi(u)$ and $\theta_i = \theta/2$, first note from (6) and (5) that $A_i + A_j = \pi(\bar{u})(1 - \frac{\theta}{2}) - \frac{\theta}{2}\pi^m = \frac{A_j}{x_i}$, such that $x_i = \frac{A_j}{A_i + A_j}$. i) then follows

directly. For ii)-iii), simply substitute x_i into (3) to give $\bar{\pi}_i = \frac{\theta}{2}\pi^m + A_j$. To derive iv), it is sufficient to substitute x_i into (2) to give $\alpha_i = 1 - \frac{A_i + A_j}{(1-\theta)\pi^m}$, and into (4) to obtain $F_i(u) = \frac{[\pi^m - \pi(u)] + [A_i + A_j]}{(1-\theta)\pi(u)}$. Lower expected utility offers then follow because an increase in either advertising cost produces a decrease in α_i and increase in $F_i(u)$ in the sense of first-order stochastic dominance, for both $i = a, b$. \square

Proof of Proposition 4. Given $A_i = A$ and $\theta_i = \theta/2$, first note from (6) and (5) that $(1 - \tau_i)x_j = (1 - \tau_j)x_i$, such that $x_i = \frac{1-\tau_i}{2-\tau_i-\tau_j}$. i) then follows directly. For ii) and iii), simply substitute x_i into (3) to give $\bar{\pi}_i = \frac{\theta}{2}(1 - \tau_i)\pi^m + \frac{1-\tau_i}{1-\tau_j}A$. To derive v), it is sufficient to substitute x_i into (2) to give $\alpha_i = \alpha_j = 1 - \frac{2-\tau_i-\tau_j}{(1-\tau_i)(1-\tau_j)} \frac{A}{(1-\theta)\pi^m}$, and into (4) to obtain $F_i(u) = F_j(u) = \frac{\frac{\theta}{2}[\pi^m - \pi(u)] + \frac{2-\tau_i-\tau_j}{(1-\tau_i)(1-\tau_j)}A}{(1-\theta)\pi(u)}$. Lower expected utility offers then follow because an increase in either τ_i or τ_j produces a reduction in α_i and an increase in $F_i(u)$ in the sense of first-order stochastic dominance, for both $i = a, b$. \square

Proof of Proposition 5. The proof follows using some similar steps to Proposition 1 after defining the set of advertising firms as $K^* = \{1, \dots, k^*\}$, and noting that any firm $l \notin K^*$ with $\alpha_l = 0$ has $x_l = 0$ under our restrictions such that $\Pi_{j \neq i}(1 - \alpha_j) = \Pi_{j \neq i \in K^*}(1 - \alpha_j)$. First, we require each firm $i \in K^*$ to be indifferent between not advertising and advertising a utility slightly higher than u^m such that $\pi_i^m[\theta_i + (1 - \theta)x_i\Pi_{j \neq i}(1 - \alpha_j)] = \pi_i^m[\theta_i + (1 - \theta)\Pi_{j \neq i}(1 - \alpha_j)] - A_i$. This implies that the probability that all firms $j \neq i$ do not advertise equals:

$$\Pi_{j \neq i}(1 - \alpha_j) = \frac{A_i}{(1 - x_i)(1 - \theta)\pi_i^m}. \quad (10)$$

After plugging this back into the previous equation, we gain $\bar{\pi}_i = \theta_i\pi_i^m + \frac{x_i}{1-x_i}A_i$ for each $i \in K^*$. The same expression also applies to those firms not in $j \neq K^*$ that do not advertise because these firms have $x_j = 0$ under our restrictions.

To ensure a common upper utility bound for each advertising firm, \bar{u} , we then require $\bar{\pi}_i = (1 - \theta_{-i})\pi_i(\bar{u}) - A_i$ for each $i \in K^*$. This provides an expression for x_i for each such firm, (11), which when summed over $i = 1, \dots, k^*$ and set equal to 1, also provides (12). When combined, these provide $k^* + 1$ equations to solve for $k^* + 1$ unknowns, $\{x_i, \dots, x_{k^*}\}$ and \bar{u} .

$$x_i = 1 - \frac{A_i}{\pi_i(\bar{u})(1 - \theta_{-i}) - \theta_i\pi_i^m} \geq 0, \quad (11)$$

$$\sum_{i=1}^{k^*} \left[1 - \frac{A_i}{\pi_i(\bar{u})(1 - \theta_{-i}) - \theta_i \pi_i^m} \right] = 1, \quad (12)$$

Consider the fully asymmetric case first. The LHS of (12) is decreasing in \bar{u} , and thus reaches its maximum at $\bar{u} = u^m$, and the value is $\bar{I}_k = \sum_{i=1}^k \left[1 - \frac{A_i}{(1-\theta)\pi_i^m} \right]$. It is easy to see that \bar{I}_k is increasing in k . The minimum is reached at $\bar{u} = \tilde{u}_k$, because \bar{u} cannot grow further and ensure that all x_i are non-negative. The value of the LHS at $\bar{u} = \tilde{u}_k$ is $\underline{I}_k \equiv \sum_{i=1}^k \left[1 - \frac{A_i}{\pi_i(\tilde{u}_k)(1-\theta_{-i}) - \theta_i \pi_i^m} \right]$. It is easy to see that \underline{I}_k is increasing in k . This is because the last element in the sum is 0 (by definition of \tilde{u}_k), so when k is reduced by 1, we get a new sum where all the $k - 1$ remaining elements decrease because of replacing \tilde{u}_k with a higher \tilde{u}_{k-1} . Finally, note that $\underline{I}_k < \bar{I}_{k-1}$ which follows from the fact that both sums have equal number of non-zero elements, and those in \underline{I}_k are evaluated at $\bar{u} = \tilde{u}_k$ where each element of the same is lower than at $\bar{u} = u^m$ where they are evaluated in \bar{I}_{k-1} .

For the equilibrium we need the largest k such that $\underline{I}_k \leq 1 < \bar{I}_k$. If this inequality holds for k , then we can find $\bar{u} \in (u^m, \tilde{u}_k]$ such that (12) holds. Since \underline{I}_k is decreasing in k , then (12) will also hold for any $k' < k$, however k' cannot be the equilibrium number of advertising firms because then firm k will be able to profitably deviate from not advertising by advertising \bar{u} . This follows from the fact that $\underline{I}_k < 1$ and so the solution to (12) for $k' < k$ has the property that $\bar{u} > \tilde{u}_k$. Thus the equilibrium k^* should be such that either $\underline{I}_{k^*} \leq 1 < \underline{I}_{k^*+1}$ for $k^* < n$,²¹ or $\underline{I}_n \leq 1 < \bar{I}_n$. These give all the conditions in (12).

As for the quasi-symmetric case, the only possibility is $k^* = n$ because there is no way to exclude any of the firms from advertising without giving it a strict incentive to deviate and advertise. Thus the only requirement of that case is $\sum_{i=1}^n \frac{A_i}{h_i(\bar{u}_n)} > (n - 1) \geq \sum_{i=1}^n \frac{A_i}{(1-\theta)\pi_i^m}$ where the first part is trivially satisfied because $\sum_{i=1}^n \frac{A_i}{h_i(\bar{u}_n)} = n$.

All that now remains is to derive the unique equilibrium advertising probabilities and utility distributions for firms $i \in K^*$. This can be done using similar steps to Proposition 5. To derive the advertising probabilities, plug (11) into (10), such that $\Pi_{j \neq i}(1 - \alpha_j) = \Pi_{j \neq i \in K^*}(1 - \alpha_j) = \gamma_i(u^m)$ for $i = 1, \dots, k^*$, where $\gamma_i(u) = \frac{\pi_i(\bar{u})(1-\theta_{-i}) - \theta_i \pi_i(u)}{(1-\theta)\pi_i(u)} \leq 1$. Then by multiplying each of these k^* equations together, we get $\Pi_{i=1}^{k^*} [\Pi_{j \neq i \in K^*}(1 - \alpha_j)] = \Pi_{i=1}^{k^*} \gamma_i(u^m)$. On simplification, this equals $\Pi_{i=1}^{k^*} (1 - \alpha_i)^{k^*-1} = \Pi_{i=1}^{k^*} \gamma_i(u^m)$ such that $\Pi_{i=1}^{k^*} (1 - \alpha_i) = [\Pi_{i=1}^{k^*} \gamma_i(u^m)]^{\frac{1}{k^*-1}}$. By now taking (10) and multiplying both sides by $1 - \alpha_i$ we get $\Pi_{j=1}^{k^*} (1 - \alpha_j) = (1 - \alpha_i) \gamma_i(u^m)$, which after substitution provides a unique solution, $\alpha_i = 1 - \frac{[\Pi_{j=1}^{k^*} \gamma_j(u^m)]^{\frac{1}{k^*-1}}}{\gamma_i(u^m)}$. Similar steps can be they used to derive the unique

²¹where \bar{I}_{k^*} follows from $\underline{I}_k < \bar{I}_{k-1}$.

utility distributions, $F_i(u) = \frac{[\prod_{j=1}^{k^*} \gamma_j(u)]^{\frac{1}{k^*-1}}}{\gamma_i(u)}$, where $F_i(u^m) = 1 - \alpha_i$ and $F_i(\bar{u}) = 1$ as required.

In order to guarantee that the equilibrium exists, we need that $F'_i(u) > 0$. It is trivially satisfied for $k^* = 2$ because $F_i(u) = \gamma_j(u)$ and $\gamma'_j(u) > 0$. In general this is not the case, and thus one needs to verify this condition for a particular application. \square

Proof of Corollary 1. Every term in the sum of the LHS of (12) is between 0 and 1. As $A \rightarrow 0$, if no denominator gets to zero, the LHS converges to k^* . Given that \tilde{u}_i are all different from each other, and given that all firms that advertise should have $\tilde{u}_i \geq \bar{u}$, the only way to make LHS equal $k^* - 1$ is to set $k^* = 2$ and let \bar{u} to converge to \tilde{u}_2 . Thus, when A is sufficiently small, only two firms, 1 and 2, can advertise.

On the other extreme, if it is the case that $\sum_{i=1}^n \frac{1}{\pi_i^m} > \frac{n-1}{\pi_n^m}$, then $A \leq (1 - \theta)\pi_n^m$ can be found such that (12) has a solution. This is because at $\bar{u} = u^m$ we have the LHS of (12) equal to $n - \sum_{i=1}^n \frac{A}{(1-\theta)\pi_i^m}$, so as $A \rightarrow \frac{(n-1)(1-\theta)}{\sum_{i=1}^n \frac{1}{\pi_i^m}}$, the solution to (12) converges to u^m . Given that $\sum_{i=1}^n \frac{1}{\pi_i^m} > \frac{n-1}{\pi_n^m}$, we have $\frac{(n-1)(1-\theta)}{\sum_{i=1}^n \frac{1}{\pi_i^m}} < (1 - \theta)\pi_n^m$ so that as $A \rightarrow \frac{(n-1)(1-\theta)}{\sum_{i=1}^n \frac{1}{\pi_i^m}}$, $\tilde{u}_n \gg u^m$, thus equilibrium \bar{u} can be found. \square

Proof of Corollary 2. The claims about A_i and θ_i follow directly from the definition of \tilde{u}_i and the characterization of advertising firms. For the multiplicative profit function the inverse function is $\pi_i^{-1}(x) = \pi^{-1}(\frac{x}{1-\tau_i})$ so $\tilde{u}_i = \pi^{-1}\left(\frac{(1-\tau_i)\pi^m + A_i}{1-\tau_i}\right) = \pi^{-1}\left(\frac{\pi^m + A_i/(1-\tau_i)}{1-\theta_i}\right)$. It is now clear that \tilde{u}_i is decreasing in τ_i . \square

Appendix B - Further Technical Equilibrium Details

B1. Market Equilibrium with Asymmetric Firms and $A_a = A_b \rightarrow 0$

When the firms are asymmetric but $A_a = A_b = A \rightarrow 0$, the equilibrium depends upon $\pi_a^{-1}\left(\frac{\theta_a \pi_a^m}{1-\theta_b}\right) \geq \pi_b^{-1}\left(\frac{\theta_b \pi_b^m}{1-\theta_a}\right)$. Without loss of generality, suppose $\pi_i^{-1}\left(\frac{\theta_i \pi_i^m}{1-\theta_j}\right) < \pi_j^{-1}\left(\frac{\theta_j \pi_j^m}{1-\theta_i}\right)$ such that $\pi_i(\bar{u})(1 - \theta_j) - \theta_i \pi_i^m < \pi_j(\bar{u})(1 - \theta_i) - \theta_j \pi_j^m$. From (5), it must be that $x_j > x_i$. Moreover, with the additional use of (6), for \bar{u} to exist and for x_i and x_j to be well defined, it must be that $\pi_i(\bar{u})(1 - \theta_j) - \theta_i \pi_i^m \rightarrow 0$ such that $x_i \rightarrow 0$, $x_j \rightarrow 1$, and $\bar{u} \rightarrow \pi_i^{-1}\left(\frac{\theta_i \pi_i^m}{1-\theta_j}\right)$. Given this, we know $\lim_{A \rightarrow 0} \bar{\pi}_i = \theta_i \pi_i^m$ and $\lim_{A \rightarrow 0} \bar{\pi}_j = \lim_{A \rightarrow 0} (1 - \theta_i) \pi_j(\bar{u}) = (1 - \theta_i) \pi_j\left(\pi_i^{-1}\left(\frac{\theta_i \pi_i^m}{1-\theta_j}\right)\right) > \theta_j \pi_j^m$. Further, from (4), we know $\lim_{A \rightarrow 0} F_i(u) = \lim_{A \rightarrow 0} \frac{\bar{\pi}_j - \theta_j \pi_j(u)}{(1-\theta) \pi_j(u)}$

and $\lim_{A \rightarrow 0} F_j(u) = \lim_{A \rightarrow 0} \frac{\bar{\pi}_i - \theta_i \pi_i(u)}{(1-\theta)\pi_i(u)}$. Finally, from (2), $\alpha_j \rightarrow 1$, while firm i advertises with probability $\lim_{A \rightarrow 0} \alpha_i = 1 - \frac{\bar{\pi}_j - \theta_j \pi_j^m}{(1-\theta)\pi_j^m} \in (0, 1)$.

Unit Demand Example: Suppose $u_i = V_i - p_i$ and $\pi_i(u_i) = V_i - c_i - u_i$, where $u_i^m = 0$, and $\pi_i^m = V_i$. By using the results above, the equilibrium then depends upon $(1 - \theta_a)(V_a - c_a) - (1 - \theta_b)(V_b - c_b) \leq 0$. For instance, when this is negative, $x_a \rightarrow 0$ and $x_b \rightarrow 1$, such that $\bar{\pi}_a \rightarrow \theta_a(V_a - c_a)$, and $\bar{\pi}_b \rightarrow (1 - \theta_a)((V_b - c_b) - \bar{u})$, where $\bar{u} \rightarrow \left(\frac{(1-\theta)(V_a - c_a)}{1-\theta_b} \right)$. By then denoting $\Delta V = V_a - V_b$, and noting that $F_a(u_b) = Pr(u_a \leq u_b) = 1 - F_a(p_b + \Delta V)$ and $F_b(u_a) = 1 - F_b(p_a - \Delta V)$, it follows that $F_a(p) = 1 - \left[\frac{\bar{\pi}_b - \theta_b(p - \Delta V - c_b)}{(1-\theta)(p - \Delta V - c_b)} \right] = 1 + \frac{\theta_b}{1-\theta} - \frac{(1-\theta_a)(\theta_a(V_a - c_a) + (1-\theta_b)(c_a - c_b - \Delta V))}{(1-\theta_b)(1-\theta)(p - \Delta V - c_b)}$ on $[V_a - \bar{u}, V_a)$ and $F_b(p) = 1 - \left[\frac{\bar{\pi}_a - \theta_a(p + \Delta V - c_a)}{(1-\theta)(p + \Delta V - c_a)} \right] = 1 - \left[\frac{\theta_a(V_b - p)}{(1-\theta)(p + \Delta V - c_a)} \right]$ on $[V_b - \bar{u}, V_b)$, where $\alpha_b \rightarrow 1$ but where firm a refrains from advertising with probability $1 - \alpha_a = 1 - F_a(V_a) \in (0, 1)$.

B2. Equilibrium with Downward-Sloping Demand

Given $\mathbf{p}_i = \{p_{i1}, \dots, p_{iK_i}\}$, the individual consumer product demand functions at firm i can be permitted to be interrelated, as summarized by the demand vector $\mathbf{q}_i(\mathbf{p}_i) = \{q_{i1}(\mathbf{p}_i), \dots, q_{iK_i}(\mathbf{p}_i)\}$. One can then write $u_i = S(\mathbf{p}_i, \mathbf{q}_i^*(\mathbf{p}_i))$, where $S(\mathbf{p}_i, \mathbf{q}_i(\mathbf{p}_i))$ denotes the indirect utility available at firm i for a given level of demand, and where $\mathbf{q}_i^*(\mathbf{p}_i)$ denotes a consumer's optimal demand vector at firm i , $\mathbf{q}_i^*(\mathbf{p}_i) = \arg\max_{\mathbf{q}_i} S(\cdot)$. It then follows that $\pi_i(\mathbf{p}_i) = \mathbf{q}_i^*(\mathbf{p}_i)'(\mathbf{p}_i - \mathbf{c}_i)$, where $\mathbf{c}_i = \{c_{i1}, \dots, c_{iK_i}\}$.

Under monopoly, firm i would set a vector of monopoly prices, $\mathbf{p}_i^m = \arg\max_{\mathbf{p}_i} \pi_i(\mathbf{p}_i)$, with $u_i^m = S(\mathbf{p}_i^m, \mathbf{q}_i^*(\mathbf{p}_i^m))$ and $\pi_i(u_i^m) \equiv \pi_i(\mathbf{p}_i^m)$. Hence, for Assumption U to hold with $u_a^m = u_b^m$, we restrict attention to cases with $\pi_a(u) = \pi_b(u) = \pi(u)$.

Under suitable demand assumptions, there can exist a unique efficient price vector that maximizes a firm's profits subject to the constraint of supplying a given utility draw u , such that $\mathbf{p}^*(u) = \arg\max_{\mathbf{p}} \pi(\mathbf{p})$ subject to $S(\mathbf{p}, \mathbf{q}^*(\mathbf{p})) = u$, with resulting profits per consumer, $\pi(u) \equiv \pi(\mathbf{p}^*(u))$.²² It then follows that $\bar{\pi}_i = \theta_i \pi(\mathbf{p}^m) + \frac{x_i}{1-x_i} A_i$, $\alpha_i = 1 - \frac{A_j}{x_i(1-\theta)\pi(\mathbf{p}^m)}$, and $F_i(u) = \frac{x_i \theta_j [\pi(\mathbf{p}^m) - \pi(\mathbf{p}^*(u))] + A_j}{x_i(1-\theta)\pi(\mathbf{p}^*(u))}$, where $\bar{\mathbf{p}} = \mathbf{p}^*(u^m) = \mathbf{p}^m$ and $\underline{\mathbf{p}} = \mathbf{p}^*(\bar{u})$, and where x_i and \bar{u} follow from amended versions of (5) and (6).

To consider how our framework then reproduces the standard clearinghouse equilibrium, suppose that the market is symmetric. It then follows that $x_i = 0.5$. Further, let $K = 1$ such that $\mathbf{p} \equiv p$, $\bar{\pi} = \frac{\theta}{2} \pi(p^m) + A$, $1 - \alpha = \left(\frac{2A}{(1-\theta)\pi(p^m)} \right)$, and $\pi(\bar{u}) = \left(\frac{\theta \pi(p^m) + 4A}{(2-\theta)} \right)$.

²²This constrained pricing decision can be thought of as Ramsey problem. Individual prices can be hard to fully characterize, but with additional restrictions, firms can be shown to optimally use lower prices on products that are more price-elastic and complementary to other products. See Armstrong and Vickers (2001) and Simester (1997) for more discussion.

We can then use $F(p) = 1 - F(u)$ to find the price distribution (conditional on advertising) $F_A(p) \equiv \frac{1-F(u)}{\alpha}$ which equals $\frac{1}{\alpha} \left[1 - \left(\frac{\theta[\pi(p^m) - \pi(p)] + 4A}{2(1-\theta)\pi(p)} \right) \right]$ with $\underline{p} = \pi^{-1} \left(\frac{\theta\pi(p^m) + 4A}{(2-\theta)} \right)$ and $\bar{p} = p^m$. Finally, to consider how our framework reproduces the equilibrium of Simester (1997), suppose the market is symmetric with $K \geq 1$ and $A \rightarrow 0$, and let all marginal costs equal zero. One can then replicate the equilibrium using $x_i = 0.5$ under the additional restriction that loyalists and shoppers share common demand functions.

Appendix C: Relaxing the Visit Assumptions

In this appendix, we provide further details on how the duopoly model can be generalized to allow for multiple visits and non-negligible visit costs. Suppose that the cost of visiting any first firm is s_1 and the cost of visiting any second firm is s_2 . The main model implicitly assumed that $s_1 = 0$ and $s_2 = \infty$. However, we now show that our equilibrium remains for any $s_2 > 0$ provided that the costs of any first visit are not too large, $s_1 \in [0, u^m]$, and that shoppers can only purchase from a single firm. Similar (but more tedious) arguments can also be made for the n -firm model.

First, suppose $s_1 \in [0, u^m]$ but maintain $s_2 = \infty$. Beyond the case of $s_1 = 0$, this now also permits cases where the costs of a first visit are strictly positive provided $u^m > 0$ as consistent with downward-sloping demand and linear prices. In particular, provided that the monopoly utility is larger than s_1 , consumers will still be willing to make a first visit and the equilibrium will remain unchanged.

Now suppose that $s_2 > 0$ subject to a persistent assumption of one-stop shopping, such that a consumer cannot buy from more than one firm. By definition, the behavior of the loyal consumers will remain unchanged. Therefore, to demonstrate that our equilibrium remains robust, we need to show that shoppers will endogenously refrain from making a second visit. Initially suppose that firms keep playing their original equilibrium strategies and that a shopper receives $h \in 0, 1, 2$ adverts. Given $s_2 > 0$ and the assumption of one-stop shopping, the gains from any second visit will always be strictly negative for all h . In particular, if $h \geq 1$, then a shopper will first visit the firm with the highest advertised utility, $u^* < u^m$, and any offer from a second visit would necessarily provide $u < u^*$. Alternatively, if $h = 0$, then both firms will offer u^m , and so any second visit is unable to offer a gain in utility. Now suppose that firms can deviate from their original equilibrium strategies. To see that the logic still holds, note that only the behavior of any non-advertising firms is relevant and that such firms are unable to influence any second visit decisions due to their inability to communicate or commit to any $u < u^m$. Hence,

firms' advertising and utility incentives remain unchanged and the original equilibrium still applies.

References

- ALLEN, J., R. CLARK, AND J.-F. HOUDE (2014): "The Effect of Mergers in Search Markets: Evidence from the Canadian Mortgage Industry," *American Economic Review*, 104, 3365–3396.
- AN, Y., M. R. BAYE, Y. HU, J. MORGAN, AND M. SHUM (2015): "Identification and Estimation of Online Price Competition With an Unknown Number of Firms," *Journal of Applied Econometrics*, *forthcoming*.
- ANDERSON, S., A. BAIK, AND N. LARSON (2015): "Personalized pricing and advertising: An asymmetric equilibrium analysis," *Games and Economic Behavior*, 92, 53–73.
- ARMSTRONG, M. AND J. VICKERS (2001): "Competitive Price Discrimination," *RAND Journal of Economics*, 32, 579–605.
- ARNOLD, M., C. LI, C. SALIBA, AND L. ZHANG (2011): "Asymmetric Market Shares, Advertising And Pricing: Equilibrium With An Information Gatekeeper," *Journal of Industrial Economics*, 59, 63–84.
- ARNOLD, M. A. AND L. ZHANG (2014): "The Unique Equilibrium in a Model of Sales with Costly Advertising," *Economics Letters*, 124, 457–460.
- BARRON, J. M., B. A. TAYLOR, AND J. R. UMBECK (2004): "Number of Sellers, Average Prices, and Price Dispersion," *International Journal of Industrial Organization*, 22, 1041–1066.
- BAYE, M. R. AND C. G. DE VRIES (1992): "Mixed strategy trade equilibria," *Canadian Journal of Economics*, 281–293.
- BAYE, M. R., D. KOVENOCK, AND C. G. D. VRIES (1992): "It Takes Two to Tango: Equilibria in a Model of Sales," *Games and Economic Behavior*, 4, 493–510.
- BAYE, M. R. AND J. MORGAN (2001): "Information Gatekeepers on the Internet and the Competitiveness of Homogeneous Product Markets," *American Economic Review*, 91, 454–474.

- (2009): “Brand and price advertising in online markets,” *Management Science*, 55, 1139–1151.
- BAYE, M. R., J. MORGAN, AND P. SCHOLTEN (2004a): “Price Dispersion in the Small and in the Large: Evidence from an Internet Price Comparison Site,” *Journal of Industrial Economics*, 52, 463–496.
- (2004b): “Temporal price dispersion: Evidence from an online consumer electronics market,” *Journal of Interactive Marketing*, 18, 101–115.
- (2006): *Information, Search, and Price Dispersion*, Amsterdam: Elsevier, vol. 1, chap. 6.
- BROWN, J. R. AND A. GOOLSBEE (2002): “Does the Internet Make Markets More Competitive? Evidence from the Life Insurance Industry,” *Journal of Political Economy*, 110, 481–507.
- BURDETT, K. AND K. L. JUDD (1983): “Equilibrium Price Dispersion,” *Econometrica*, 51, 955–969.
- CHANDRA, A. AND M. TAPPATA (2011): “Consumer Search and Dynamic Price Dispersion: An Application to Gasoline Markets,” *The RAND Journal of Economics*, 42, 681–704.
- CHIOVEANU, I. (2008): “Advertising, brand loyalty and pricing,” *Games and Economic Behavior*, 64, 68–80.
- DUBOVIK, A. AND M. C. JANSSEN (2012): “Oligopolistic Competition in Price and Quality,” *Games and Economic Behavior*, 75, 120 – 138.
- GIULIETTI, M., M. WATERSON, AND M. WILDENBEEST (2014): “Estimation of Search Frictions in the British Electricity Market,” *The Journal of Industrial Economics*, 62, 555–590.
- GU, Y. AND T. WENZEL (2014): “Strategic obfuscation and consumer protection policy,” *The Journal of Industrial Economics*, 62, 632–660.
- HONG, P., R. P. MCAFEE, AND A. NAYYAR (2002): “Equilibrium Price Dispersion with Consumer Inventories,” *Journal of Economic Theory*, 105, 503–517.

- HOSKEN, D. AND D. REIFFEN (2004): “Patterns of Retail Price Variation,” *RAND Journal of Economics*, 35, 128–146.
- (2007): “Pricing Behavior of Multiproduct Retailers,” *The B.E. Journal of Theoretical Economics*, 7, 1354–1354.
- IRELAND, N. J. (1993): “The provision of information in a Bertrand oligopoly,” *The Journal of Industrial Economics*, 61–76.
- JANSSEN, M. C. AND J. L. MORAGA-GONZÁLEZ (2004): “Strategic pricing, consumer search and the number of firms,” *The Review of Economic Studies*, 71, 1089–1118.
- JING, B. (2007): “Product differentiation under imperfect information: When does offering a lower quality pay?” *Quantitative Marketing and Economics*, 5, 35–61.
- KOCAS, C. AND T. KIYAK (2006): “Theory and Evidence on Pricing by Asymmetric Oligopolies,” *International Journal of Industrial Organization*, 24, 83–105.
- LACH, S. (2002): “Existence and Persistence of Price Dispersion: An Empirical Analysis,” *The Review of Economics and Statistics*, 84, 433–444.
- LEWIS, M. (2008): “Price dispersion and competition with differentiated sellers,” *The Journal of Industrial Economics*, 56, 654–678.
- MORAGA-GONZÁLEZ, J. L., Z. SÁNDOR, AND M. R. WILDENBEESE (2013): “Semi-Nonparametric Estimation Of Consumer Search Costs,” *Journal of Applied Econometrics*, 28, 1205–1223.
- NAKAMURA, E. AND J. STEINSSON (2008): “Five Facts about Prices: A Reevaluation of Menu Cost Models,” *The Quarterly Journal of Economics*, 123, 1415–1464.
- NARASIMHAN, C. (1988): “Competitive Promotional Strategies,” *The Journal of Business*, 61, 427–449.
- PENNERSTORFER, D., P. SCHMIDT-DENGLER, N. SCHUTZ, C. WEISS, AND B. YONTCHEVA (2015): “Information and Price Dispersion: Theory and Evidence,” *Working paper*.
- RAJU, J. S. (1995): “Theoretical models of sales promotions: Contributions, limitations, and a future research agenda,” *European Journal of Operational Research*, 85, 1–17.

- ROBERT, J. AND D. O. STAHL (1993): “Informative Price Advertising in a Sequential Search Model,” *Econometrica*, 61, 657–86.
- ROY, S. (2000): “Strategic segmentation of a market,” *International Journal of Industrial Organization*, 18, 1279–1290.
- SCHMIDT, R. C. (2013): “Price competition and innovation in markets with brand loyalty,” *Journal of Economics*, 1–27.
- SHELEGIA, S. (2012): “Asymmetric Marginal Costs in Search Models,” *Economics Letters*, 116, 551–553.
- SHERMAN, J. AND A. WEISS (2015): “Search and Price Dispersion at an Open-Air Market,” *Working paper*.
- SIMESTER, D. (1997): “Optimal Promotion Strategies: A Demand-Sided Characterization,” *Management Science*, 43, pp. 251–256.
- SORENSEN, A. T. (2000): “Equilibrium Price Dispersion in Retail Markets for Prescription Drugs,” *Journal of Political Economy*, 108, 833–862.
- STAHL, D. O. (1989): “Oligopolistic Pricing with Sequential Consumer Search,” *American Economic Review*, 79, 700–712.
- STEENKAMP, J.-B. E., V. R. NIJS, D. M. HANSSENS, AND M. G. DEKIMPE (2005): “Competitive reactions to advertising and promotion attacks,” *Marketing Science*, 24, 35–54.
- VARIAN, H. R. (1980): “A Model of Sales,” *American Economic Review*, 70, 651–659.
- WILDENBEEST, M. R. (2011): “An empirical model of search with vertically differentiated products,” *The RAND Journal of Economics*, 42, 729–757.