

Asymmetric Price Effects of Competition*

Saul Lach[†]
The Hebrew University and CEPR

José L. Moraga-González[‡]
Vrije Universiteit Amsterdam

December 22, 2016

Abstract

In markets where price dispersion is prevalent the relevant question is not what happens to *the price* when the number of firms changes but, instead, what happens to *the whole distribution of equilibrium prices*. Using data from the gasoline market in the Netherlands, we find, first, that markets with a given number of competitors have price distributions that first-order stochastically dominate the corresponding price distributions in markets with one more firm. Second, the competitive response varies along the price distribution and is stronger at prices in the medium to upper part of the distribution. Finally, simulations of the consumer gains from competition reveal that they depend on how well informed consumers are and would be larger for relatively attentive consumers. A generalisation of Varian's (1980) model allowing for richer heterogeneity in consumer price information along the lines of Burdett and Judd's (1983) model can account for these empirical patterns.

Keywords: price dispersion, number of competitors, distribution of price information

JEL Classification: D43, D83, L13

*We thank the Editor, Alessandro Gavazza, and two anonymous referees, from whom we received numerous comments that helped us improve the paper. We are also indebted to Marco Haan, Pim Heijnen and Adriaan Soetevent for providing us with the dataset for this study and for answering our many questions about it. We also thank Guillermo Caruana, Riemer Faber, Chaim Fersthman, David Genesove, Alex Gerhskov, Dietmar Harhoff, Corine Hoeben, Maarten Janssen, Gerard Llobet, Jan de Loecker, Benny Moldovanu, Marielle Non, Carlos Noton, Vaiva Petrikaite, Giuseppe Ragusa, Zsolt Sandor, Dirk Stelder, Matthijs Wildenbeest, Mariano Tappata and Shlomo Yitzhaki for their comments. The paper has also benefited from seminar participants at the ASSET Meetings 2010, Ben Gurion University, CEMFI, 10th CEPR Conference on Applied Industrial Organization, CERGE-EI, Instituto Di Tella, 24th Annual Congress of the EEA, 37th Annual Congress of the EARIE, The Hebrew University, London School of Economics, Tel-Aviv University, University of Mannheim, University of Zurich, 3rd Workshop on Search and Switching Costs and The World Bank. Boaz Abramson, Yaron Aronshtan, Danny Bahar, Eli Berglas and Sarit Weisburd provided excellent research assistance.

[†]E-mail: <saul.lach@mail.huji.ac.il>. Lach thanks the Wolfson Family Charitable Trust for its financial support.

[‡]E-mail: <j.l.moragagonzalez@vu.nl>. Moraga is also affiliated with the University of Groningen, The Tinbergen Institute, CEPR, CESifo and the Public-Private Sector Research Center (IESE, Barcelona). Moraga gratefully acknowledges financial support from the Marie Curie Excellence Grant MEXT-CT-2006-042471.

1 Introduction

Economists have long studied the relationship between the number of firms and prices. Standard Cournot and Bertrand oligopoly models assume consumers are perfectly informed about all prices in the market and predict that an increase in the number of competitors lowers the equilibrium price. These standard models formalize the widely accepted view in economics that more competition, as measured by an increase in the number of competitors, lowers prices and benefits all consumers.

Alternative, and more realistic, models depart from the assumption that all consumers have the same information and, in the absence of the possibility to price discriminate, describe equilibria characterized by price distributions.¹ In such markets where price dispersion is prevalent the relevant question is no longer what happens to *the* price when the number of firms changes but, instead, what happens to *the whole distribution of equilibrium prices*. Moreover, when market equilibria are characterised by price distributions, increased competition may affect consumers differentially, depending on whether they are more or less alert to price information. These observations lead us to ask: How do the different percentiles of the price distribution change as competition intensifies? Do all consumers gain from competition? Which consumers gain more, those well informed about prices or those poorly informed?

This paper studies how the distributions of posted and paid prices change with the number of competitors in a real-world market where price dispersion is common, namely, the gasoline market in The Netherlands. We use station-level daily prices of Euro 95 gasoline obtained from 3,195 gas stations for the 8-month period March 3, 2006–October 31, 2006. In order to measure the number of competitors of a gas station i , we define the (local) market for a gas station i as the geographical area in a circle with radius r centered at the location of gas station i . To estimate the causal effect of competition on prices we cannot use panel data methods like in Gerardi and Shapiro (2009) because we do not have variation over time in the number of competitors active in a given market. Our identification strategy, instead, is based on price comparisons between markets *within the same municipality* with different numbers of competitors under the assumption that the variation in other characteristics affecting *both* prices and competition across markets in the same municipality is negligible.

¹The information consumers have about prices varies across them for exogenous and/or endogenous reasons. *Exogenous* reasons include the involuntary exposure to price information when consumers, for example, walk around, talk to friends and neighbours, and/or interact in social networks. *Endogenous* reasons include active consumer search as well as advertising and promotional activities. For a recent survey of models that rationalize price dispersion, see Baye et al. (2006).

After providing evidence that gas stations frequently change their prices in a way that is consistent with mixed strategies, we first show that all the estimated effects of an increase in the number of competitors on the deciles of the price distribution are negative and statistically significant. This implies that markets with a given number of gas stations have price distributions that first-order stochastically dominate the corresponding price distributions in markets with one more firm. Second, the estimated effects vary significantly with the decile. The estimated coefficients for the medium to upper deciles are larger (i.e., more negative) than those for the lower deciles. That is, the competitive response varies along the price distribution and is stronger at prices in the medium to upper part of the distribution. Third, the estimated impact of an additional gas station becomes weaker as the market area increases and the typical gas station faces more competitors.

These empirical findings suggest that all consumers do benefit from an increase in the number of firms. However, because not all the deciles respond in the same way to an increase in competition, it is possible that consumer gains are asymmetrically distributed across the consumer population. Based on the idea that the nature of price dispersion stems from the existence of differentially informed consumers, which is natural in our data because frequent price changes make it more likely that consumers have imperfect price information, we perform simulations of the consumers gains from competition and show that the largest gains would accrue to consumers who are relatively well informed about prices. Arguably, if lower to medium income consumers have more incentives to pay attention to the posted prices while they drive around or commute, our results suggest that competition is not only efficiency enhancing but also serves to redistribute rents towards those income segments.

Empirical research of markets with price dispersion has usually proceeded by estimating the impact of competition on the mean and variance of prices. Using data from US gasoline markets, Marvel (1976), Barron et al. (2004) and Lewis (2008) find evidence that a higher number of sellers is associated with a lower mean price and price dispersion. For the Austrian gasoline market, Pennerstorfer et al. (2015) also find that more competition results in lower average prices, though they relate competition to the share of informed consumers in a market, which they measure as the share of long-distance commuters. Price dispersion exhibits an inverted-U relationship with respect to the proxy for the share of informed consumers. Our empirical results are broadly consistent with these studies. Research on the effects of competition on prices in other retail markets includes Baye et al. (2004), Borenstein and Rose (1994) and Gerardi and Shapiro (2009). Baye et al. (2004) study price

dispersion in online markets and conclude that a higher number of vendors lowers price dispersion. In a seminal paper about the airline industry Borenstein and Rose (1994) found greater amounts of price dispersion in routes with higher numbers of competitors. However, Gerardi and Shapiro (2009) found the opposite in a similar study and argued that the reason for the disparity between their results and those of Borenstein and Rose had to do with the estimation method. Our results are consistent with the findings of Baye et al. (2004) and Gerardi and Shapiro (2009).

While studying how the mean and variance of prices respond to higher levels of competitiveness is certainly informative, focusing only on these two statistics is much too narrow to uncover the clear patterns we observe in our data, namely, that more competition causes the entire distribution to shift to the left and, moreover, that price dispersion falls because the medium and top deciles of the price distribution fall more than the bottom deciles do. Our broader approach also provides the interesting welfare implication that more competition benefits consumers unequally. The only paper we are aware of making similar points is a study of a bank merger in the Canadian mortgage market by Allen et al. (2014). Unlike us, they observe the prices actually paid by consumers. They find evidence that markets with more competitors have lower but more dispersed negotiated mortgage rates. Like our paper, they also look at the distributional effects of the merger and they find that the merger has a bigger impact on the relatively well informed consumers. We only observe the posted prices but our simulations of the paid prices are consistent with their finding.

The classical price dispersion models of Varian's (1980) and Rosenthal's (1980) fail to account for the empirical finding that price distributions shift to the left as the number of competitors increases. In fact, their models predict, instead, that an increase in the number of competitors raises the average price.² This distinct prediction has sometimes been used to discriminate among possible explanations of observed price dispersion and Varian-type models have been dismissed as a plausible explanation on this ground (e.g., Barron et al., 2004; Hosken et al., 2008; Haynes and Thompson, 2008). One way to accommodate our empirical results is to allow for endogenous consumer search, see e.g. Dana (1994), Burdett and Judd (1983) and Janssen and Moraga-González (2004). However, we do not believe that much search takes place in gasoline markets because consumers hardly deviate from their commuting path in order to shop for lower prices. Instead, in order to accommodate our empirical findings, and

²Specifically, in the famous shopper/non-shopper formulation of Varian (1980), an increase in the number of firms shifts the "upper" portion of the price distribution to the right while the "lower" portion shifts to the left, and this occurs in such a manner that the mean price goes up (Varian himself did not prove this latter result; for a proof see Janssen and Moraga-González, 2004, and Morgan et al., 2006). Rosenthal's (1980) model predicts that an increase in the number of competitors shifts the entire price distribution to the right.

following Armstrong, Vickers and Zhou (2009), we propose to modify Varian's (1980) model by allowing for more general forms of consumer price information heterogeneity in order to capture consumer variation in driving and commuting patterns. We study such a generalised model where we postulate that if a market has, say, 4 firms then the share of consumers who know exactly 1, 2, 3 and 4 prices is positive.³ Under the weak assumption that an increase in the number of competitors lowers the share of consumers who observe one price only and increases the average number of prices observed in the market, we find plausible conditions under which the price distribution shifts to the left as competition increases. Examples in which our theoretical model fits relatively well the observed effects in the data are when the distribution of price information in the market follows a (discrete) uniform or a (truncated) binomial distribution. In such cases, an increase in the number of firms shifts the entire distribution to the left in such a way that consumers who are more alert to the posted prices derive greater utility gains from an increase in competition.⁴

We believe the message of this paper goes beyond the present application to the gasoline market in the Netherlands. Since imperfect price information is prevalent in many markets (telecommunications, health, gas, electricity, etc.), the price effects of competition-enhancing policies (industry deregulation, trade liberalization, merger control, etc.) might not be as straightforward as those implied by standard models. Moreover, since increased competition can potentially have unequal effects among consumers, distributional issues become a central part of the welfare assessment of these policies. This advocates the importance of taking a broader view where the interaction between competition and consumer policy is taken into consideration (Armstrong, 2008; Waterson, 2003; Wu and Perloff, 2007).

The remainder of the paper is organized as follows. Our empirical analysis of the gasoline market in the Netherlands is in Section 2. In Section 3 we propose a model, inspired by Varian's (1980) model of sales, of the distribution of prices in an oligopolistic market where consumers differ in the amount of prices they are exposed to. Conditions are derived under which prices can increase or decrease as the number of competitors goes up. Proofs are relegated to the Appendix. The paper closes with Section 4, where we offer some concluding remarks.

³In Varian's shopper and non-shopper formulation the assumption is that consumers observe either 1 or 4 prices but never 2 or 3.

⁴More heterogeneity in consumer access to price information is also the central tenet in Moraga-González et al. (2016) and Nermuth et al. (2013), which focus on how increased search affects price levels, price dispersion and consumer welfare.

2 Empirical findings from the retail gasoline market

We examine gasoline prices in the Netherlands to provide motivation for the main theoretical results of this paper. First, we provide evidence consistent with the use of mixed strategies by gas stations and then examine how competition, measured by the number of gas stations, affects the whole distribution of prices in the market. Our focus is on the approach to analyzing the price effects of competition in markets characterized by price dispersion in equilibrium and heterogeneous consumer price information, rather than on the particular results of this case-study. When competition affects the distribution of prices in an asymmetric way, the welfare of differentially-informed consumers will also be affected in an uneven way because they will pay prices drawn from different (minimum) price distributions. Understanding the effects of changes in competition throughout the price distribution can therefore inform policy better than estimates of its impact on just the mean and variance of prices, which has typically been the focus of the literature.

We use daily gasoline (Euro 95) prices from a large sample of gas stations in the Netherlands. The price data were obtained from Athlon Car Lease Nederland B.V., the largest private car leasing company in the Netherlands with over 129,000 cars as of the end of 2008 (www.athloncarlease.com).⁵ Athlon leases cars to private companies and these cars are used by their employees for work purposes. The typical contract between Athlon and its lessees stipulates that Athlon pays for the gasoline consumed as well as for car maintenance, repair, insurance, etc. Fleet-cards are used to make these payments, and it is the information in these cards that enables Athlon to learn the prices paid by each of the drivers of its cars. Athlon's lessees do not get special discounts so the prices reported by Athlon are the actual prices posted by the gas stations.⁶

Prices were obtained from 3,195 gas stations for the 8-month period March 3, 2006–October 31, 2006. This number does not include gas stations located along highways since these stations face a different market environment than stations located within cities. The sample covers 80 percent of the universe of non-highway gas stations in the Netherlands.⁷ Because the price information arrives

⁵The data used in this paper are part of the data collected and analyzed by Soetevent *et al.* (2014). We are indebted to them for providing us with the data. Soetevent *et al.* (2014) study whether ownership changes in highway gas stations originating from a government program of auctions and divestitures enhances competition. For further details on the data collection, see their Appendix B.

⁶Compared to using price data from credit card transactions, using price data from fleet cards does not have the problem that the prices are likely to be biased towards the lower part of the price distribution.

⁷The 3,195 distinct gas stations represent 3,260 legal entities since some gas stations changed ownership during the sample period. We treat these cases as a single station. The Dutch Competition Authority reports that 4,319 legal entities were registered in the Netherlands in 2005, including 244 located in highways. Our data cover, therefore, 80 percent of non-highway gas stations in the Netherlands $\left(\frac{3260}{4319-244}\right)$. Note that this is a very high coverage because the figure reported by the Dutch Competition Authority includes stations unavailable to the public (owned by garages),

directly from the lessees, not all stations are sampled every day which results in an unbalanced panel data of gas stations.

The following institutional details about the data allow us to view the sample as random draws from the distribution of prices. First, the price data come from drivers of lease cars that use the car on behalf of their employer and hence do not have to pay for their own gasoline bill (it is part of the leasing contract). It is then reasonable to assume that these drivers have no incentives to search for the gas stations offering the lowest prices. This is important because, otherwise, our sample would have been a sample of gas stations charging low prices. On the other hand, these consumers may be choosing to buy from gas stations offering better service values and therefore charging higher prices. The fact that our sample covers 80 percent of the universe of gas stations should allay these concerns. It could still occur that Athlon's lessees patronize a selected subset of the stations and this will be reflected in a disproportionately larger number of price observations for these stations. In one of our robustness checks we will control for this possible selection by restricting the sample to a balanced panel of gas stations. Second, all gas stations in the Netherlands are self-service and therefore there is a single price for gas in each station.⁸ Finally, the available information leads us to believe that the extent to which various prices in a given market are set by a single firm (because of joint-ownership) and/or reflect collusive agreements is minor.⁹

There are 306,436 station-day observations on Euro95 gasoline prices. For illustration, Figure 1 displays the density function of prices in the sample (solid line). The lowest price in our sample is 102 cents while the highest price is 167 cents. The mean and median price of gasoline is 138.9 and 140 cents, respectively, and the standard deviation is 7.15 cents. This variation is mainly driven by changes in prices over time. Removing the temporal variation by using residuals from a regression of observed prices on a set of day and brand dummies accounts for a great deal of the observed variation (residual prices have a standard deviation of 2.62). Their density (after adding the mean

stations in construction and sites that only sell diesel.

⁸For a paper identifying price discrimination based on self- and full-service in the gasoline market see Shephard (1991).

⁹Although we do not have information on the gas stations' owners, according to the Dutch Competition Authority, about 62 percent of the gas stations are owned and operated by independent dealers (NMa, 2006). The remaining stations belong to the main oil producers: BP, Esso, Shell, Texaco and Total. But even among these branded stations, most are dealer-operated. For example, Shell serves fewer than 15 percent of the gas stations and about 2/3 of the Shell-branded gas stations are operated by dealers who are free to set their own prices. This suggests that joint-ownership of gas stations is not such a prevalent phenomenon as one may be led to believe from casual observation (although we have no data on joint-ownership of gas stations by independent owners). An exception was the highway market, where most gas stations, 63 percent, were owned and operated by the large oil producers (NMa, 2006). However, starting in 2002, the Dutch government has forced divestitures of highway stations in order to increase competition. This is one of the main reasons for not including stations located along highways in our sample.

price) is the dash curve in Figure 1.

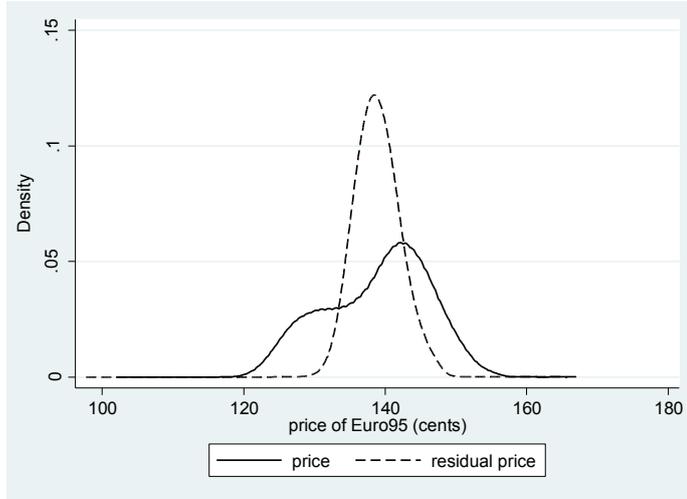


Figure 1: Density of Euro95 prices

Not surprisingly, there is dispersion in gasoline prices and, although not very large, it has some economic significance. As an illustration, the largest difference between the highest and lowest prices offered by gas stations within the same municipality and during the same day was 38 cents; this occurred on May 1 in the municipality of Langedijk and on June 3 in the municipality of Boxtel. This implies that a consumer filling a 50-liter tank at the lowest-priced station instead of at the highest-priced station would have saved 19 euros.

We define the (local) market for a gas station i as the geographic area within a circle of radius r centered at the location of gas station i . The number of local markets is therefore the same as the number of gas stations. This approach has been taken in other studies of the gasoline market.¹⁰ Soetevent et al. (2014), using the same data, found that rivals' price effects are significant up to $2km$, and strongest at distances of $1km - 2km$. We therefore consider two values for the radius r : $1km$ and $2km$. We use the list of gas stations (and their geographic locations) to compute the number of firms N_i operating in market i (including gas station i).¹¹ Thus, station i is active in a market facing $N_i - 1$ rivals. Since we have information on 80 percent of all the stations in the Netherlands, N_i may be underestimated in some cases.

We restrict the original sample of 3,195 gas stations or local markets (306,436 station-day obser-

¹⁰See, for example, Shepard (1991), Hastings (2004), Lewis (2008), Chandra and Tapatta (2011), and Soetevent et al. (2014). For an alternative definition of markets based on commuting paths see Houde (2012) and Pennerstorfer et al. (2015). We do not adopt this approach because the main commuting routes in the Netherlands are highways and we exclude gas stations along highways from the analysis.

¹¹Markets overlap and N_i is also the number of markets in which station i is present. Obviously, N_i depends on the radius r used to define the local market but we omit this dependency from the notation.

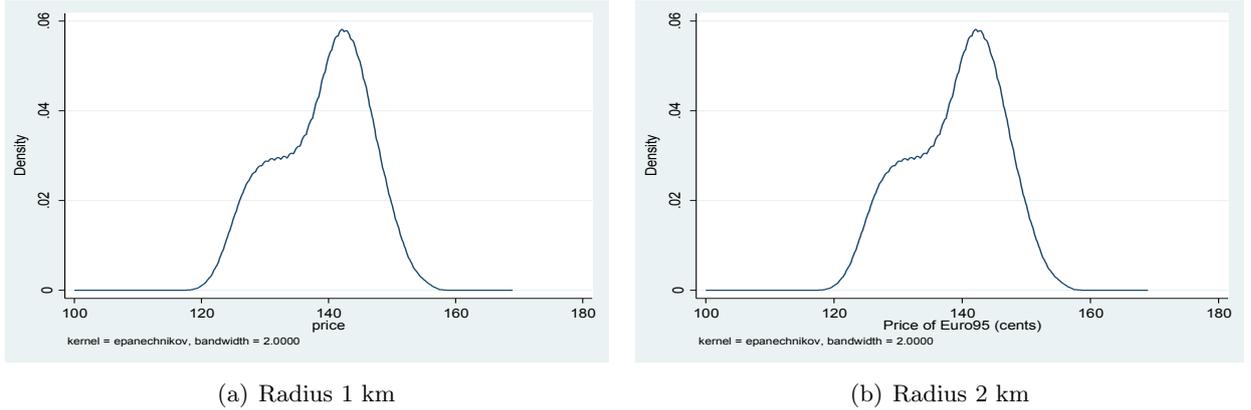


Figure 2: Density of Euro95 prices in markets with different radii (excluding monopolies and large municipalities)

vations) in two ways. First, we omit markets where there is only one gas station, i.e., $N_i = 1$. We do this because the price dispersion models mentioned in the Introduction do not typically apply to monopolies (nor does the model we present in Section 3). This reduces the number of gas stations to 1,465 when r is $1km$ and to 2,443 when r is $2km$.¹² The much larger reduction when $r = 1km$ reflects the fact that the number of markets with $N_i = 1$ decreases considerably with market size (determined by the radius r). Second, we restrict the sample to small municipalities in order to minimize within-municipality variation in unobservables affecting both prices and the number of gas stations (more on this in Section 2.3). Specifically, we use gas stations located in municipalities in the lower two thirds of the municipality area distribution, i.e., municipalities with an area less than $90.4 km^2$. This further reduces the number of gas stations to 796 (81,095 observations) and 1,344 (141,119 observations) when r is $1km$ and $2km$, respectively. Figure 2 indicates, however, that the density of prices in each of these subsamples is essentially identical to the one in the original sample. In fact, the mean, median and standard deviation of prices are almost identical. In any case, as a robustness check, we will examine the sensitivity of our results to the inclusion of monopolies in the sample and to variations in the threshold defining small municipalities (Section 2.3.1).

Table 1 shows the distribution of the number of gas stations N_i across markets. Note that increasing the radius from $1km$ to $2km$ considerably lowers the number and share of markets with two competitors. The median number of stations across markets varies between 2 and 4.

In our sample, gas stations are usually not located in front of each other. This is more likely to occur along highways but highway stations are excluded from our analysis. In fact, the shortest

¹²The number of observations is reduced to 144,102 and 245,084, respectively.

Table 1: Distribution of the number of gas stations (N)

	Radius:1 km				Radius:2 km			
	Stations	Observations	Stations	Observations	Stations	Observations	Stations	Observations
Mean	2,45	2,46	3,92	3,95				
Median	2	2	4	4				
Std. Dev.	0,73	0,74	1,76	1,74				
N	Frequency	Percent	Frequency	Percent	Frequency	Percent	Frequency	Percent
2	518	65,1	52.682	64,96	332	24,7	32.674	23,2
3	215	27,0	22.095	27,25	316	23,5	33.905	24,0
4	49	6,2	4.703	5,8	279	20,8	30.458	21,6
5	10	1,3	1.128	1,39	173	12,9	18.730	13,3
6	2	0,3	277	0,34	121	9,0	12.493	8,9
7	1	0,1	145	0,18	68	5,1	7.193	5,1
8	1	0,1	65	0,08	34	2,5	3.368	2,4
9	0	0	0	0	10	0,7	1.137	0,8
10	0	0	0	0	7	0,5	616	0,4
11	0	0	0	0	4	0,3	545	0,4
Total	796	100	81.095	100	1.344	100	141.119	100

Excluding markets with monopolies and located in large cities. See text for details.

distance between stations is 30 meters and only 24 pairs of gas stations are within 100 meters of each other. When $r = 1km$, the mean (median) distance between station i and its closest rival is 560 (592) meters, and 75 percent of the 796 stations are at least 344 meters away from its closest competitor. Similarly, when $r = 2km$, the mean (median) distance between station i and its closest rival is 895 (876) meters, and 75 percent of the 1,344 stations are at least 504 meters away from its closest competitor.

The panel data set is unbalanced. For each gas station i we do not observe its prices in every day of the 8 month sample period. For $r = 1km$, the median number of days when prices are observed is 116 days per station (the mean is 102 days). The maximum number of price observations per station is 148 days, and the minimum is just one day. A quarter of the 796 gas stations has less than 70 daily observations, while another quarter has more than 139 daily observations. The panel structure is similar when $r = 2km$.¹³

The unbalanced nature of the panel also implies that we do not observe the prices of the $N_i - 1$ competitors of station i in every day. For each station and for every day when its prices are observed we indicate whether the $N_i - 1$ prices of i 's competitors are also observed. For each station, we average this indicator over time to get the percentage of days in which we observe complete market price data. The mean and median of these percentages across stations is 0.55 and 0.59, respectively, when $r = 1km$, and 0.41 and 0.37, respectively, when $r = 2km$. This implies that, for the average gas station, we will be able to rank prices on 55 (41) percent of the days when its prices are observed

¹³The median number of days for which we observe prices in a station is 120 days (the mean is 105 days). The maximum number of price observations per station is 148 days, and the minimum is just one day. A quarter of the 1,344 gas stations has less than 79 daily observations, while another quarter has more than 140 daily observations.

when $r = 1km$ ($2km$).

2.1 Mixed strategies

Price dispersion models prescribe firms to use mixed strategies in equilibrium. Hosken *et al.* (2008), Lewis (2008), Chandra and Tappata (2011), among others, present evidence on price dispersion and behavior consistent with the use of mixed strategies for different samples of gasoline stations located in the US, while Pennerstorfer *et al.* (2015) do so for the case of Austria.

The use of mixed strategies implies that we should not observe gas stations always selling at high prices or always selling at low prices.¹⁴ In this subsection we provide supportive evidence of this prediction by checking whether the gas stations in our sample vary their relative position in the cross-sectional distribution of prices over time. Following Chandra and Tappata (2011) and Pennerstorfer *et al.* (2015) we compute a measure of rank reversals to capture variation in the stations' ranking over time. For a station i and a rival j in i 's market such that i 's price is equal to or higher than j 's price in over 50 percent of the days where both prices are observed, we count the number of times j 's price is higher than i 's price, i.e., we count the occurrences of rank reversals. Formally, we define

$$rr_{ij} = \begin{cases} \frac{1}{|T_{ij}|} \sum_{t \in T_{ij}} I(p_{jt} > p_{it}) & \text{if } \frac{1}{|T_{ij}|} \sum_{t \in T_{ij}} I(p_{it} \geq p_{jt}) > 0.5 \\ \text{missing} & \text{otherwise} \end{cases}$$

where T_{ij} is the set of days for which we observe both prices, $|T_{ij}|$ is the cardinality of this set, and $I(\cdot)$ is the indicator function.

The mean value of rr_{ij} over all pairs (i, j) for which rr_{ij} is computed is 0.098 (0.105) and the fraction of non-zero cases is above 67 (79) percent when $r = 1km$ ($2km$). This means that a station that usually charges the higher price also posts the lowest price about 10 percent of the time (out of a possible maximum of 50 percent). These values are about equal to those reported by Pennerstorfer *et al.* (2015) and slightly lower than those in Chandra and Tappata (2011).

2.2 Stochastic dominance

In this subsection we estimate, and compare, the cumulative distribution functions (CDF) of prices for markets with different levels of competition. We estimate the CDF of prices by pooling all the

¹⁴Bunching of prices is not much of an issue in our data. If we define price bunching as an event in which all prices charged by the N stations in the market are equal, then this occurs in 19.7 and 8.2 percent of the market-day observations when markets are defined by radius 1km and 2km, respectively. If we define price bunching less restrictively as a situation where at least one pair of prices in the same market are the same, then this occurs in 26.7 and 21.7 percent of the pairs when markets are defined by radius 1km and 2km, respectively. Moreover, the existence of bunching per-se is not evidence against the use of mixed strategies as mixed strategy models based on differentially informed consumers (along the lines of the one presented in Section 3) that allow for firm asymmetries deliver mixed strategy equilibria with price distributions having atoms (see Baye *et al.*, 1992).

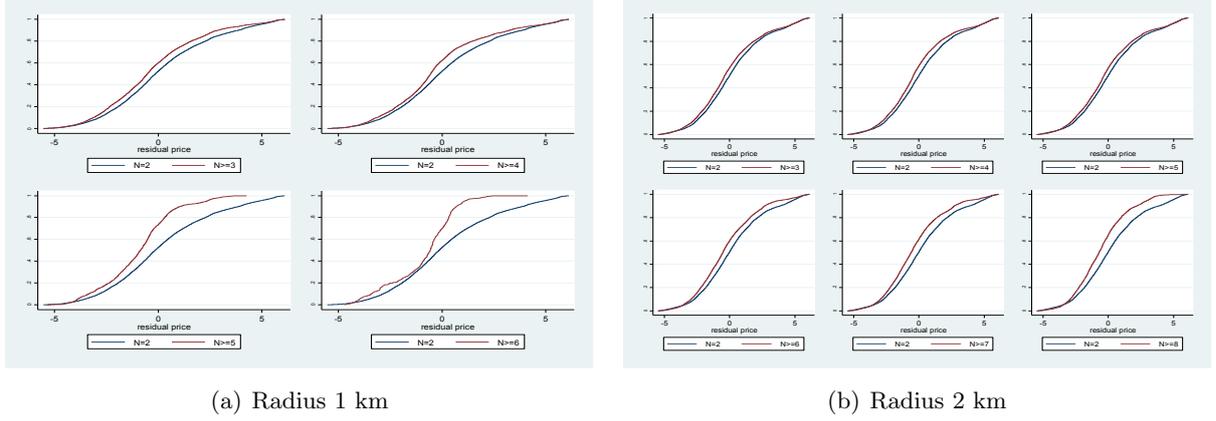


Figure 3: Empirical cumulative distribution functions

price observations for all gas stations (markets) conditional on N . In order to pool observations from different gas stations in different days we net out the effects of time and brand by using residuals from a regression of observed prices on a set of day and brand dummies. In each panel of Figure 3 we contrast the empirical CDF of residual prices in markets with $N = 2$ to that in markets with $N \geq n$ for $n \in [3, 6]$ when $r = 1km$ and $n \in [3, 8]$ when $r = 2km$.¹⁵

It is quite clear that the price distribution when $N = 2$ lies below the price distribution in markets with more than 2 competitors. Moreover, it appears that the larger N is, the further away the distributions are. The same is true for other contrasts (e.g., $N = 3$ and $N \geq n$ for $n \geq 4$). Table 2 presents results of the Kolmogorov-Smirnov test of the null hypotheses that both CDFs are the same against the one-sided hypothesis that the CDF when $N = 2$ is smaller than the CDF when $N \geq n$.¹⁶

Table 2: Kolmogorov-Smirnov test for equality of distributions

	Radius: 1km		Radius: 2km	
	statistic	p-value	statistic	p-value
$N > 3$	-0,0777	0	-0,0679	0
$N > 4$	-0,1025	0	-0,0843	0
$N > 5$	-0,2353	0	-0,0651	0
$N > 6$	-0,2676	0	-0,0967	0
$N > 7$	-	-	-0,1102	0
$N > 8$	-	-	-0,1521	0

Entries are the Kolmogorov-Smirnov statistic and p-value for testing the null hypothesis that the CDF when $N=2$ and when $N \geq n$ are equal against the alternative hypothesis that the CDF when $N=2$ is smaller than the CDF when $N \geq n$. Asymptotic p-values are presented.

¹⁵We restrict N to $N \leq 6$ ($N \leq 8$) because the number of observations declines drastically when $N \geq 7$ ($N \geq 9$) for $r = 1km$ ($2km$) (see Table 1). Furthermore, the CDFs were estimated after trimming the data by excluding the bottom and top 1 percent of the observations in order to avoid extreme values of residual prices.

¹⁶This test should be taken as suggestive only because it assumes independent observations which is not likely to be the case here, even though day effects have been removed and there are gaps in the time-series data.

The null hypothesis is always rejected in favor of the alternative that the price distribution with 2 competitors is smaller than the price distribution with more than n competitors. For example, the largest difference between both distributions when there are more than 4 competitors is 0.1025 and 0.0843 for r equal $1km$ and $2km$, respectively, and these differences are significantly different from zero. Conversely, the null hypotheses cannot be rejected at conventional levels of significance when the alternative is that the distribution with $N \geq n$ is smaller than the price distribution with 2 competitors (results not shown). Thus, the distribution of prices when $N = 2$ stochastically dominates the price distribution when $N \geq n$.¹⁷

2.3 Quantile regressions

In this subsection we pursue a complementary parametric approach to estimate the effect of competition on the distribution of prices. Specifically, we estimate how the number of competitors affects percentiles of the price distribution using quantile regression methods.

The canonical quantile regression model can be expressed as follows (e.g., Koenker, 2004):

$$p_{it} = x_{it}\beta_\tau + u_{\tau,it} \quad q_\tau(p_{it}|x_{it}) = x_{it}\beta_\tau \quad (1)$$

where p_{it} is the price of gas station i at time (day) t , x_{it} is a vector of regressors, β_τ is a corresponding vector of parameters to be estimated, $u_{\tau,it}$ is an error term and $q_\tau(p_{it}|x_{it})$ is the τ^{th} quantile of p_{it} conditional on the regressors x_{it} .

The estimator of β_τ – from the τ^{th} quantile regression – solves

$$\underset{\beta}{Min} \sum_{i,t: p_{it} \geq x_{it}\beta} \tau |p_{it} - x_{it}\beta| + \sum_{i,t: p_{it} < x_{it}\beta} (1 - \tau) |p_{it} - x_{it}\beta|$$

or, equivalently,

$$\underset{\beta}{Min} \sum_i \sum_{t \in T_i} \rho_\tau(p_{it} - x_{it}\beta)$$

where $\rho_\tau(u) = (\tau - I(u < 0))u$ is the “check” or asymmetric absolute loss function, T_i is the set of days in which prices of station i are observed, and the outer sum is over all the stations in the sample.

¹⁷Similar significant results obtain when using actual prices instead of residual prices but the differences between the CDFs are of lower magnitude, and when using residual prices obtained from a regression of prices on brand dummies and a long distributive lag of wholesale gasoline prices instead of day dummies.

In contrast to the least squares estimator which estimates the effect of regressors at the conditional mean of the dependent variable, quantile regression estimates such effects at different quantiles of its distribution determined by the choice of τ . The coefficient $\beta_{\tau k}$ can be interpreted as the partial effect of a particular regressor x_{itk} on the conditional quantile of p , $\frac{\partial q_{\tau}(p_{it}|x_{it})}{\partial x_{itk}}$, i.e., the change in the price corresponding to the τ^{th} conditional quantile due to a small change in x_{itk} . By choosing different values of τ , quantile regression allows us to track the effect of changes in x_{itk} on different parts of the distribution of the dependent variable.

In our application, the focus is on the effect of the number of gas stations in station i 's market, N_i , on different quantiles of the price distribution. As in the conventional OLS framework, the causal interpretation of the partial effect of N_i depends upon N_i being exogenous.¹⁸ We therefore want to control for factors that affect both prices and the number of gas stations in the market. As an example, consumers may differ in their willingness to pay for gasoline and in their shopping behavior because of reasons as varied as differences in income, commercial activity and spatial-related factors. Usually these demand factors affect both prices and profitability (and therefore the number of gas stations) in the same direction. If we do not control for these factors an endogeneity problem arises and the estimated effect of N on prices will likely be upward biased, i.e., biased towards zero if the causal effect is negative.

Time-invariant unobserved factors are typically controlled for by using fixed effects for the gas stations. In the context of quantile regressions this is not as straightforward to do as in mean regression because the nonlinearity of the quantile equation does not allow for the “differencing-out” of the fixed effects. In any case, we cannot implement such an approach in our data since N_i is constant over the sample period. The length of the sample period, 8 months, is too short to observe much entry and, in fact, we do not have any records of entry episodes in our data. We therefore address the endogeneity concern by adding fixed effects for the municipality where the gas station is located. Thus, to a first approximation, identification of the effect of N_i on the quantiles of the price distribution is obtained by relating the variation in prices to the variation in the number of competitors among gas stations *within the same municipality*.¹⁹ The assumption is that, within municipalities, the variation in the number of competitors is not related to unobserved factors affecting prices. For example, in Nuth, a small municipality (16,000 inhabitants) in the southern part of the Netherlands, there are four gas

¹⁸The implicit assumption in (1) is $q_{\tau}(u_{\tau,it}|x_{it}) = 0$.

¹⁹Most of the price variation observed in Figure 1 is within municipalities: the between-municipality variation accounts for only 6.9 percent of the total variation in price.

stations with different number of competitors (using the $2km$ market definition): two stations have $N = 3$ while the other two have $N = 4$ and $N = 5$. The assumption is that unobserved demand factors do not vary enough within Nuth to explain the variation in N across the four gas stations.²⁰ This variation can therefore be treated as exogenous.

The assumption of negligible variation in local (within-municipality) demand factors is more likely to hold the smaller the size of the municipality, and it is precisely for this reason that we focus our empirical analysis on “small” municipalities. The 3,195 gas stations in the original sample were located in 440 municipalities, the majority of which are small geographic units both in terms of area and population: the mean (median) area is 92 km^2 (63.4 km^2) and the mean (median) population is 37,000 (22,900). There are, however, some very large municipalities in the sample (e.g., Amsterdam, Rotterdam). We therefore use the lower 2/3 part of the municipality area distribution, i.e., municipalities having an area less than 90.4 km^2 . As shown in Table 1, after excluding monopolies and stations in large municipalities we are left with 796 (1,344) gas stations located in 184 (246) municipalities when markets are defined by radius $1km$ ($2km$). These municipalities are less spread-out than in the original sample but have, on average, the same population size.²¹ It should be remarked that there is variation in N within municipalities: 40 percent of the 184 municipalities used when $r = 1km$ exhibit variation in N , and this number increases to 87 percent when markets are defined by $r = 2km$.

We also control for other demand-side factors such as brand preferences and within week temporal variation in shopping patterns by using brand and day-of-the-week dummy variables. Supply-side factors are less of a problem in our context because much of the operating costs are determined by wages and the wholesale price of gasoline. Both of these are common to most gas stations and it is quite likely that they do not vary within municipalities.

We use daily price data over the 8 month sample period and therefore need to account for possible variation in the distribution of prices over time. Prices vary over time mainly because of day-to-day variation in the wholesale price of gasoline. We use a distributive lag of the daily spot price of gasoline from the Amsterdam-Rotterdam-Antwerp (ARA) spot market to account for these variations.

²⁰As explained in Section 2, the markets overlap and therefore each station in Nuth appears as a competitor in many markets. Moreover, the station having four competitors, i.e., $N = 5$, has one competitor from a neighboring municipality (Schinnen). In such cases, we need to assume that neighboring municipalities are sufficiently similar to each other which is a reasonable assumption in the Netherlands.

²¹When $r = 1km$, the mean (median) area of the 184 municipalities in the sample is 43.8km^2 (38.4km^2) and the mean (median) population is 38,800 (27,000). When $r = 2km$, the mean (median) area of the 246 municipalities in the sample is 43.5km^2 (38.6km^2) and their mean (median) population is 33,400 (25,500).

The empirical specification of the quantile function in (1) is

$$p_{it} = q_{\tau}(p_{it}|x_{it}) + u_{\tau,it} = \sum_{\ell=0}^L \omega_{\ell\tau} c_{t-\ell} + \delta_{\tau} \ln N_i + z_{it} \delta_{z\tau} + \mu_{m(i)\tau} + u_{\tau,it} \quad (2)$$

where $c_{t-\ell}$ is the (lagged) wholesale gasoline price at date $t - \ell$, z_{it} is a vector of day-of-the-week and (time-invariant) brand dummies, $\mu_{m(i)\tau}$ is a municipality fixed effect measured at the municipality m where i is located, and $u_{\tau,it}$ is an error term assumed to be quantile independent of N_i conditional on the other covariates. We use 14 lags ($L = 14$) to capture the serial correlation in prices.²²

The logarithm approximation is a parsimonious way of capturing a decaying effect of an increase in N which, as in standard oligopoly models, is a sensible assumption.

We estimate the parameter β_{τ} by quantile regressions methods for $\tau \in \{0.1, 0.2, \dots, 0.9\}$.²³ Table 3 shows the estimates of δ_{τ} only for both market definitions. The estimated coefficients are proportional to the response of the τ^{th} price quantile to an increase in the number of firms from N to $N + 1$, $\delta_{\tau} \ln [1 + 1/N]$.²⁴ Standard errors are robust to heteroskedasticity. Ideally, one would also like to correct for any spatial correlation left among stations in the same market (after accounting for the municipality fixed effect) and any serial correlation left (after accounting for 14 lags of the wholesale price) in the unobserved error. This, however, is challenging because markets overlap so that simple clustering methods do not work as argued by Pennerstorfer et al. (2015). Pennerstorfer et al. (2015) propose an alternative approach, applicable to mean regressions, that allows for spatial autocorrelation and find that standard errors increase by a small amount. In our case, the heteroskedasticity-robust standard errors are so low that even if more robust variances were computed it is unlikely that our conclusions will be altered.²⁵

A few results are worth noticing. First, most estimated effects are negative and statistically significant (the effects are not statistically different from zero in two cases). This implies that markets with fewer gas stations have price distributions that first-order stochastically dominate the corresponding price distributions with more firms. This is consistent with the graphical evidence presented in Section 2.2. Second, the estimated effects vary with the quantile τ . For the 1 km markets, the estimated coefficients clearly increase (in absolute value) with the quantile. In this case,

²²We experimented with longer lags such as $L = 30$ but the estimates of interest remained unchanged.

²³We do not estimate more extreme quantiles because the asymptotic distribution of estimators of extreme quantiles is non-standard (see Chernozhukov, 2005).

²⁴The factor $\ln [1 + 1/N]$ starts at 0.69 when $N = 1$ and declines rapidly with N , being equal to 0.15 when $N = 6$.

²⁵We thank an anonymous referee for pointing this out. As a robustness check, in the Online Appendix we implement a different procedure that produces standard errors that control for spatial and serial correlation. As expected, the resulting standard errors are larger than those in the results of Table 3 but the estimates remain significantly different from zero (in 8 of the 9 deciles).

Table 3: Quantile regression estimates

τ	Estimate of δ	
	Radius: 1km	Radius: 2km
0,1	-0,382*** 0,060	-0,056** 0,028
0,2	-0,345*** 0,047	-0,145*** 0,024
0,3	-0,355*** 0,045	-0,183*** 0,023
0,4	-0,390*** 0,047	-0,211*** 0,023
0,5	-0,386*** 0,048	-0,184*** 0,023
0,6	-0,473*** 0,051	-0,154*** 0,025
0,7	-0,640*** 0,058	-0,129*** 0,028
0,8	-0,879*** 0,072	-0,035 0,035
0,9	-1,145*** 0,095	-0,024 0,052
# municipalities	184	246
# stations	796	1.344
# observations	81.095	141.119

Note: All quantile regressions include current and 14 lags of the wholesale price of gasoline as well as brand ,day-of-the-week and municipality dummies.

Small numerals are standard errors robust to heteroskedasticity.

* p<0.10; ** p<0.05; *** p<0.01

the competitive response is strongest at the upper part of the distribution. For the 2 km markets, the effect is non-monotonic, with the strongest impact at the center of the price distribution. These results suggest that increased competition has asymmetric price effects: it affects different parts of the price distribution differently. Third, the estimated effects are smaller when the market area is larger (radius $2km$ compared to radius $1km$). This is to be expected because larger markets contain more gas stations on average and hence the impact of an additional gas station is weaker.

It is instructive to compare our findings to those obtained from the standard approach analyzing the effect of competition on the mean and dispersion of prices. In order to do this, we follow the standard practice and regress prices on the same set of regressors used in (2), compute residuals and then regress the squared residuals on those regressors. Table 4 reports the coefficient of $\ln N$ from the mean and variance regressions.

We observe that competition has a negative effect on both the mean and variance of prices.²⁶ However, the relationship between the mean/variance of prices and competition is not informative

²⁶A negative relationship between price dispersion and the number of gas stations is often found in the empirical literature (e.g., Barron et. al. (2004), Lewis (2008)), although there are exceptions such as Chandra and Tappata (2011). Borenstein and Rose (1994) also found that price dispersion in airline prices increases with competition. However, as shown by Gerardi and Shapiro (2009), this result is due to the presence of omitted variables and disappears once these are controlled for using panel data. See Baye et al. (2006) for a review of the theoretical and empirical literature.

Table 4: Effect of competition on the mean and variance of prices

	Radius: 1 km	Radius: 2 km
Mean		
Log (number of stations)	-0,562*** (0,042)	-0,160*** (0,021)
R-squared	0,912	0,900
Variance (residual squared)		
Log (number of stations)	-1,212*** (0,144)	-0,171*** (0,078)
R-squared	0,189	0,111
# stations	796	1.344
# observations	81.095	141.119

Note: All regressions include current and 14 lags of the wholesale price of gasoline as well as brand, day-of-the-week and municipality dummies. Small numerals are standard errors robust to heteroskedasticity.

about which parts of the price distribution are changing, and by how much, as competition increases. In contrast, the quantile regressions in Table 3 provide this type of information.

To summarize our results, we note that the data clearly show that increased competition shifts the entire price distribution to the left in a non-uniform or asymmetric way. Moreover, the magnitude and patterns of these competitive effects depend on the geographic size of the market.

2.3.1 Robustness checks

In this subsection, whose details can be found in the Online Appendix to this paper, we reestimate the quantile regressions using different subsamples of data to check the robustness of our conclusions. In particular, we do 5 new estimations. First, we include stations in all the municipalities in the Netherlands (not just in the “smaller” ones). Second, we keep stations in all municipalities except the very large ones by removing those in the top 10 percent of the area distribution. Third, we add monopolies to the baseline sample. Fourth, we keep only observations with $N = 2$ or $N = 3$. Finally, we use only gas stations with approximately the same number of price observations. The results are in Tables A1 and A2 in the Online Appendix. Overall, these robustness checks reinforce our conclusion that i) more competition decreases the distribution of prices, ii) the competition effect is asymmetric, and iii) the magnitude and patterns of the competitive effects depend on the market definition.

2.4 Benefits from competition

We now proceed to assess the extent to which consumers benefit from an increase in competition. Our estimates reveal that the magnitude of the effect of increased competition on prices is not very

large: an increase from 2 to 3 gas stations decreases the median price between 0.075 ($-0.184 \times \ln 1.5$) and 0.16 cents ($-0.386 \times \ln 1.5$), depending on the market definition, which is not a large effect, even relative to the small observed variation in prices across gas stations (about 7 cents).

Recall, however, that the distribution of posted prices is not the same as the distribution of prices actually paid by consumers. The latter depends on the number of prices observed by the consumer and its response to competition depends on how different parts of the equilibrium price distribution are affected by it. For example, if competition has a much stronger effect on higher than on lower prices then the expected price paid is likely to decrease more for consumers observing a small number of prices than for consumers observing a large number of prices. Thus, understanding the way the whole distribution of prices changes with competition – and not only its mean and variance – is crucial if we are interested in an assessment of the welfare implications of competition on various consumer segments.

Because we lack data on how informed consumers are, we analyze this issue by simulating the expected price paid by an hypothetical consumer who observes s prices in the market of a station i . Once we compute the expected price paid, we investigate how this price changes with N_i .

We proceed as follows. For each local market i and each day when prices are observed for all N_i gas stations we draw s prices randomly and without replacement and *store* the minimum of these s prices.²⁷ This minimum price should be seen as the simulated price paid by a consumer observing s prices. For each s , we repeat this simulation 1,000 times for each local market and day and compute the mean of these 1,000 minimum prices (by market and day). This is the estimate of the expected price paid when observing s prices in market i in day t and we denote it by \hat{y}_{its} .

We are interested in how competition affects \hat{y}_{is} for a given s and in comparing these effects across s . We therefore regress \hat{y}_{is} on $\ln(N_i)$ controlling for the same regressors used in the quantile regressions (municipality, day-of-the-week and brand dummies, as well as current and 14 lagged wholesale gasoline prices). We run a separate OLS regression for $s = 1, 2, \dots, 5$, i.e., for consumers that observe up to 5 prices. We stop at 5 because there are few markets with N above 5 (see Table 1) and hence, as s increases, the number of observations declines dramatically. In addition, we should probably not expect large differences among the prices paid by consumers who compare 5 or more prices. Results are in Table 5.

²⁷Restricting the sample to days when all competitors have observed prices reduces the number of observations. When $r = 1km$, the number of stations is slightly reduced from 796 to 774 but the number of observations declines from 81,095 to 47,087. When $r = 2km$ the number of stations is reduced from 1,344 to 1,258 but the number of observations declines from 141,119 to 60,795.

Table 5: Expected price paid and the number of gas stations

Dep. Var: Expected price paid when observing s prices					
Radius 1 km					
Number of prices observed	s=1	s = 2	s = 3	s = 4	s =5
Log (number of stations)	-1,093*** (0,049)	-1,142*** (0,052)	-0,231** (0,107)	-0,143 (0,316)	-0,804 (2,553)
Number of observations	47087	47087	12657	1830	356
R-squared	0,942	0,933	0,959	0,965	0,973
Radius 2 km					
Number of prices observed	s=1	s = 2	s = 3	s = 4	s =5
Log (number of stations)	-0,308*** (0,028)	-0,507*** (0,028)	-0,720*** (0,048)	-1,398*** (0,073)	-0,531*** (0,156)
Number of observations	60795	60795	39392	22803	10495
R-squared	0,945	0,943	0,946	0,948	0,944
Note: All regressions include current and 14 lags of the wholesale price of gasoline as well as brand ,day-of-the-week and municipality dummies. Small numerals are standard errors robust to heteroskedasticity. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$					

The estimates are all negative.²⁸ A negative coefficient means that the prices paid by consumers decrease as the number of competitors increases. For markets defined by $r = 1km$, these effects are only significant for $s \leq 3$. This is probably due to the relatively small number of observations used in the regressions and the consequent large standard errors. For both market definitions $r = 1km$ and $r = 2km$, note that consumer gains from increased competition appear to first increase and then decrease in the number of prices they observe. In general, relatively well informed consumers ($s > 1$) benefit more than poorly informed consumers ($s = 1$).

The workhorse model of price competition when consumers are differentially informed is Varian (1980). Varian's stylised shopper/non-shopper formulation of demand, however, fails to accommodate what we observe in the data because it predicts that an increase in the number of competitors increases the average price. In fact, this distinct prediction has sometimes been used to discriminate among possible explanations of observed price dispersion and Varian-type models have been dismissed as a plausible explanation on this ground (e.g., Barron et al., 2004; Hosken et al., 2008; Haynes and Thompson, 2008). Extensions of Varian-type models with endogenous consumer search can indeed generate the result that prices decrease with the number of firms. One such model has been studied in Janssen and Moraga-González (2004). However, the assumption that there is *active search* in gasoline markets is rather unrealistic because consumers hardly change their driving and commuting patterns to shop for lower prices. In the next section, therefore, we proceed to a generalization of

²⁸When $s = 1$ the mean observed price paid equals the mean posted price and therefore the estimates in column $s = 1$ should be similar to the mean regression results in Table 4. The difference could be due to the fact that the current estimates are based on a subsample of the sample used in Table 4.

Varian’s (1980) model that generates the patterns observed in our data. Inspired by Burdett and Judd’s (1983) model of non-sequential search, we propose to relax the simple shopper/non-shopper formulation of Varian by allowing for more general forms of consumer information heterogeneity.

3 A model for the distribution of prices

The market for gasoline is a good example of a homogenous goods market where price dispersion is observed.²⁹ Part of this price dispersion is persistent and due to observed factors creating gas station differentiation, such as spatial location, brand name, service provided and the additional amenities made available at the pump. However, prices change quite frequently and it is not trivial to tell which gas station is the cheapest in a given market. After controlling for observed factors, a significant amount of price dispersion remains. This points towards there being an element of strategic price mixing. One rationale for such a mixing is that, because consumers differ in their driving and commuting patterns, and because some are more attentive than others, they end up with different amounts of price information. Some consumers are informed about one price only because they happen to drive through only one gas station while on their usual route. Other consumers drive through various gas stations, pay attention to the posted prices and choose to tank up at the cheapest one. Other consumers may be less attentive and remain poorly informed even if they encounter several gas stations in their commuting path. Under these circumstances, even if gas stations are differentiated, the pricing equilibrium may be characterized by mixed strategies.³⁰

In order to learn how consumer information heterogeneity affects pricing, we use Armstrong, Vickers and Zhou’s (2009), which allows for a richer information structure than the earlier shopper/non-shopper formulation of Varian (1980) often used in the literature. This is similar to Burdett and Judd’s (1983) model of non-sequential search but they allow for endogenous search, which ends up removing a lot of the heterogeneity.³¹ We also abstract from aspects related to gas station differen-

²⁹Recent empirical studies documenting price dispersion in various markets for homogeneous products include Lach (2002) and Wildenbeest (2011) for grocery products, Hortaçsu and Syverson (2004) for mutual funds, Barron, et al. (2004), Hosken et al. (2008), and Lewis (2008) for gasoline and Baye et al. (2004) for products sold online.

³⁰Wildenbeest (2011) presents a tractable model where firms are vertically differentiated and mix in prices. In his model, prices are serially correlated because of the vertical differentiation but at the same time because of the mixing firms positions in the price rankings change from time to time. Such model fits well what we observe in gasoline price data.

³¹Varian’s (1980) model of sales is isomorphic to an all-pay auction where firms bid by cutting prices in order to win a prize consisting in the additional demand stemming from the fully informed customers (Baye et al., 1992; Moldovanu and Sela, 2001). Allowing for an arbitrary distribution of price information in the market, as we do in this paper, sets our model apart from the all-pay auction literature. First, we do have multiple heterogeneous prizes as in Barut and Kovenock (1998) but in our game a single player can win many, even all, prizes at a time. Second, since poorly informed consumers only see a few prices, a firm bidding for these consumers is only in competition with a subset of other rivals; in this sense our game is better seen as one where players participate in multiple simultaneous all-pay

tiation, but the model can easily be extended to accommodate vertical product differentiation along the lines of Wildenbeest (2011).

There are $N \geq 2$ retailers competing in prices to sell a homogeneous good to a large number L of heterogeneous consumers. At a given moment in time, a consumer wishes to purchase at most a single unit of the good.³² The maximum willingness to pay for the good of a firm is given by v . Letting c denote the unit cost of a firm, define $k \equiv v - c$.

Following Armstrong, Vickers and Zhou (2009), for $0 \leq x \leq 1$, define

$$\alpha_N(x) \equiv \sum_{s=1}^N \mu_s(N) x^s \quad (3)$$

as the *probability generating function (PGF)* for the number of prices observed by consumers. The number $\mu_s(N)$ denotes the probability of observing the prices of $s \leq N$ distinct firms, and $\sum_{s=1}^N \mu_s(N) = 1$. The PGF (3) is meant to represent the distribution of price information in the market. As mentioned above, the main source of price information heterogeneity is consumer variation in driving and commuting patterns, as well as in attentiveness to posted prices. Nevertheless, this formulation is general enough to capture other sensible sources of information such as word-of-mouth communication and social networking (Galeotti, 2010). Note that the s^{th} derivative of $\alpha_N(x)$ with respect to (wrt) x , which we will denote $\alpha_N^{(s)}(x)$, evaluated at 0 is equal to $s! \mu_s(N)$. We shall assume that the probability of observing exactly one price is strictly between 0 and 1, i.e., $\alpha_N^{(1)}(0) \equiv \mu_1(N) \in (0, 1)$.

Firms play a simultaneous-moves game. Let p_i be the price of a firm i . An individual firm i chooses its price taking the prices of the rival firms as given. There are no pure-strategy equilibria. To see this, consider the position of a firm i and suppose all its rivals were charging a price \tilde{p} , with $c \leq \tilde{p} \leq v$. Two forces affect price-setting of such firm i . First, there is a desire to *steal business* from its competitors, which pushes this firm to offer better deals than the rivals. This desire arises because the chance consumers see various price offers is strictly positive. Second, the possibility of *extracting surplus* from consumers who do not compare prices prompts firm i to offer higher prices than its rivals. This desire arises because there is a chance that consumers have no other option than buying at firm i . It is easy to see that either of these deviations destabilizes the proposed equilibrium price \tilde{p} . Therefore a single price level cannot accommodate these two incentives.³³

auctions with different number of rival players and heterogeneous prizes. To the best of our knowledge, this situation has not been studied so far.

³²This assumption is inconsequential. All our results extend to the case where consumers have downward sloping demand functions. We assume inelastic demands to ease the exposition only.

³³As in Varian (1980), when $\mu_1(N) = 1$, then all firms offering $p = v$ is a pure-strategy equilibrium. When $\mu_1(N) = 0$

Denote the mixed strategy of a firm i by a distribution of prices F_i . We shall only study symmetric equilibria, i.e., equilibria where $F_i = F$ for all $i = 1, 2, \dots, N$.³⁴ To calculate the expected profit obtained by a firm i offering the good at a price $p_i \in [c, v]$ when its rivals choose a price randomly chosen from the cumulative distribution function F , we consider the chance that firm i sells to a consumer at random. A consumer will buy from firm i if he observes the offer of firm i , which occurs with probability $s\mu_s(N)/N$, and the offer of firm i is more attractive than any other offer he receives, which happens with probability $(1 - F(p_i))^{s-1}$. The expected demand of firm i at price p_i is therefore $L \sum_{s=1}^N \frac{s\mu_s(N)}{N} (1 - F(p_i))^{s-1}$ which, using (3), is equal to $\frac{L}{N} \alpha_N^{(1)}(1 - F(p_i))$. The expected profit to firm i is

$$\Pi_i(p_i; F) = \frac{L}{N} (p_i - c) \cdot \alpha_N^{(1)}(1 - F(p_i)) \quad (4)$$

In a mixed strategy equilibrium, a firm i must be indifferent between offering any price in the support of F_i and offering the upper bound \bar{p}_i . Therefore, any price p_i in the support of F_i must satisfy $\Pi_i(p_i; F_i) = \Pi_i(\bar{p}_i; F_i)$. In symmetric equilibrium, $F_i = F$, $\bar{p}_i = \bar{p}$ and $\Pi_i = \Pi$. As a result, since $\Pi(\bar{p}; F)$ is monotonically increasing in \bar{p} , it must be the case that $\bar{p} = v$ and $\Pi(\bar{p}; F) = \frac{L}{N} k \alpha_N^{(1)}(0)$. Hence, F must solve

$$(p_i - c) \cdot \alpha_N^{(1)}(1 - F(p_i)) = k \alpha_N^{(1)}(0). \quad (5)$$

for any p in $\left[c + k \alpha_N^{(1)}(0) / \alpha_N^{(1)}(1), v \right]$, the support of F .³⁵

Unfortunately, equation (5) cannot be solved explicitly for F , except in special cases. Existence and uniqueness of an equilibrium price distribution can, however, be easily proven (see Burdett and Judd, 1983). Though it is in general impossible to obtain the equilibrium price distribution analytically, we can easily derive its inverse

$$q(\alpha_N(\tau)) = c + k \frac{\alpha_N^{(1)}(0)}{\alpha_N^{(1)}(1 - \tau)} \quad (6)$$

We note that for $\tau \in [0, 1]$, equation (6) gives the τ^{th} percentile of the equilibrium price distribution of a firm.

A close look at equilibrium condition (6) serves to make an important point: what truly matters for determining the equilibrium price distribution of a firm is not N – the number of firms – but instead, then all firms offering $p = c$ is a pure-strategy equilibrium.

³⁴Standard derivations, which can be readily adapted from, e.g., Varian (1980), show that the support of F must be a convex set and that F cannot have atoms.

³⁵Notice that the lower bound of the price distribution is always above marginal cost, which reflects the fact that firms have market power.

the distribution of information among consumers.³⁶ It is therefore changes in the distribution of information that cause more or less “competitive pressure” in the market; changes in N *per se* have no effect on prices. We summarize this simple result in:

Proposition 1 *The number of firms N affects the equilibrium price distribution only indirectly through the PGF of price information among consumers $\alpha_N(x)$.*

The model is, in fact, a model about the effect of consumers’ price information on the equilibrium price distribution. Any changes in $\alpha_N(x)$ will likely affect equilibrium prices. The focus of this paper, however, is on those changes in $\alpha_N(x)$ induced by changes in N . In this regard, we make the following assumption:

Assumption 1. *An increase in the number of firms (i) (weakly) lowers the probability consumers see one price only (i.e. $\alpha_{N+1}^{(1)}(0) \leq \alpha_N^{(1)}(0)$) and (ii) raises the (weighted) average number of prices observed in the market (i.e., $\alpha_{N+1}^{(1)}(1) > \alpha_N^{(1)}(1)$).*

Assumption 1 is a rather weak and sensible assumption. In particular, we note that it is weaker than first-order stochastic dominance (FOSD).³⁷

3.1 Equilibrium price distribution and the number of firms

Our objective is to study the relationship between prices in a market and the number of firms. Typical studies focus on the first and second moments of the price distribution. Like in our empirical study, here, we take a broader approach and examine the response of all the percentiles of the price distribution to changes in the number of competitors.

To do this, we study the impact of a change in the PGF $\alpha_N(x)$, caused by a change in N , on the (inverse) price distribution (6). The impact of an increase in N on the percentile τ of the price distribution is

$$q(\alpha_{N+1}(\tau)) - q(\alpha_N(\tau)) = k \left[\frac{\alpha_{N+1}^{(1)}(0)}{\alpha_{N+1}^{(1)}(1 - \tau)} - \frac{\alpha_N^{(1)}(0)}{\alpha_N^{(1)}(1 - \tau)} \right]. \quad (7)$$

Since k is positive, expression (7) clearly shows that the way an increase in competition affects the different percentiles of the price distribution depends on how $\alpha_N(x)$ changes into $\alpha_{N+1}(x)$.

³⁶We note that this relies on the assumption of constant returns to scale. With economies or diseconomies of scale, the decrease in quantities caused by an increase in the number of firms would have cost, and by implication, price consequences.

³⁷To be sure, for our results to hold we need that either of the inequalities $\alpha_{N+1}^{(1)}(0) \leq \alpha_N^{(1)}(0)$ and $\alpha_{N+1}^{(1)}(1) \geq \alpha_N^{(1)}(1)$ is strict. We have chosen to work with the second inequality being strict but our proofs can easily be adapted to the similar case where the first inequality is strict instead.

Proposition 2 *Suppose that the number of firms increases from N to $N + 1$ and that Assumption 1 holds. Then:*

(I) *There exists $\hat{\tau} \in (0, 1]$ such that all the percentiles of the price distribution below $\hat{\tau}$ decrease.*

(II) *All the percentiles of the price distribution will decrease if and only if*

$$\frac{\alpha_{N+1}^{(1)}(0)}{\alpha_{N+1}^{(1)}(1-\tau)} - \frac{\alpha_N^{(1)}(0)}{\alpha_N^{(1)}(1-\tau)} < 0, \text{ for all } \tau. \quad (8)$$

(III) *If*

$$\frac{\alpha_{N+1}^{(2)}(0)}{\alpha_{N+1}^{(1)}(0)} - \frac{\alpha_N^{(2)}(0)}{\alpha_N^{(1)}(0)} < 0, \quad (9)$$

*then there exists $\tilde{\tau} \in [0, 1)$ such that all the percentiles $\tau \geq \tilde{\tau}$ of the price distribution will increase.*³⁸

This result says that when an increase in the number of competitors does not increase the chance of consumers being informed about only one price and does raise the number of prices they know on average then more competitors in a market *always* results in a fall in the lower percentiles of the price distribution.

Condition (8) is a necessary and sufficient condition for the equilibrium price distribution with N firms to dominate in a FOSD sense the distribution with $N + 1$ firms. In such a case, *all* percentiles of the price distribution fall as we move from an N - to an $N + 1$ -firm market. This situation accords with the usual intuition that markets with more firms have lower prices and we remark that (8) is violated in Varian's (1980) model.³⁹ For an easier-to-interpret sufficient condition, we can state,

Corollary 1 (of Proposition 2) *A sufficient condition for all the percentiles of the price distribution to decrease is that the PGF $\alpha_N(x)$ satisfies:*

$$\frac{\alpha_{N+1}^{(s)}(0)}{\alpha_{N+1}^{(1)}(0)} - \frac{\alpha_N^{(s)}(0)}{\alpha_N^{(1)}(0)} \geq 0 \text{ for all } s = 1, 2, \dots, N$$

We note that this condition is weaker than the monotone likelihood ratio property (MLRP).⁴⁰ In other words, if the probability distribution of price information satisfies the MLRP then increased competition implies a fall in all the percentiles of the equilibrium price distribution.

The last part of Proposition 2 says that the upper percentiles of the price distribution will increase if condition (9) holds. Note that cutting prices to capture well informed consumers results in lower

³⁸In special cases, it may happen that $\alpha_{N+1}^{(2)}(0) = 0$. In those cases, we can invoke higher order derivatives. In particular, the condition would involve the lowest $s \geq 2$ for which $\alpha_{N+1}^{(s)}(0) > 0$ (see the Appendix).

³⁹In fact, in Varian's model $\mu_1(N) = \alpha_N^{(1)}(0) = \alpha_{N+1}^{(1)}(0) = \mu_1(N + 1)$ and $\mu_N(N) = 1 - \alpha_N^{(1)}(0) = 1 - \alpha_{N+1}^{(1)}(0) = \mu_{N+1}(N + 1)$. In this case, condition (8) requires $N - (N + 1)(1 - \tau) < 0$, which can never be satisfied for all $\tau \in [0, 1]$.

⁴⁰The MLRP requires that the ratio $\alpha_N^{(s)}(0)/\alpha_{N+1}^{(s)}(0)$ decreases in s (see Milgrom, 1981).

expected profits for the firms. As a result, firms try to compensate by adjusting the frequency with which they charge higher prices, thereby generating higher profits from the consumers who are less well informed about prices. As we move up in the price distribution, the prices are less and less successful at capturing well informed consumers. In effect, at the top of the price distribution, firms only care about consumers observing one or two prices because the chance of selling to other (better informed) consumers is negligible. Given this, if the probability that consumers observe one price relative to the probability that consumers observe two prices increases when we move from an N - to an $N + 1$ -firm market, then firms prefer to raise the frequency of the higher prices and so the upper percentiles of the price distribution increase.

Proposition 2 has stated conditions under which the percentiles of the price distribution increase or decrease when we move from a market with N retailers to a market with $N + 1$ retailers; however, the proposition is silent with respect to whether some percentiles increase (or decrease) more than others. As we have seen in our empirical application, this is a relevant issue because an increase in the number of competitors may be felt more in some percentiles than in others and this opens up the possibility that consumers' gains from an increase in competition be asymmetric in sign and magnitude.

To investigate this issue further, we analyze in detail a couple of examples. In Example 1 we assume that the probability of observing s prices is binomial.⁴¹ The results of this example, namely, that all percentiles of the equilibrium price distribution decrease when the number of firms increases and that those percentiles in the middle of the price distribution fall more than the lower and the upper ones is in stark contrast with the results obtained in Example 2, where we assume Varian's (1980) information structure. In the latter the middle and upper percentiles instead increase in the number of firms.

Example 1 (The truncated binomial distribution) ⁴² *The truncated binomial distribution (TBD) has PGF*

$$\alpha_N(x) = \frac{[1 - p(1 - x)]^N - (1 - p)^N}{1 - (1 - p)^N}.$$

where $p \in [0, 1]$ is the success probability of a Bernoulli experiment. The experiment consists of

⁴¹In the Online Appendix of this article, we study a third case in which the probability of observing s prices is (discrete) uniform. In that case, when $N \geq 3$, the middle percentiles of the equilibrium price distribution fall more with an increase in the number of competitors than the lower and the upper percentiles. This is similar in spirit to what we obtain assuming that the probability of observing s prices is binomial.

⁴²We refer here to the zero-truncated distribution, since our random variable –the number of prices consumers observe– has support $\{1, 2, \dots, N\}$.

observing (or not) a price and the binomial distribution gives the probability of observing s prices out of N independent trials. Note that Assumption 1 holds for the TBD. In fact,

$$\alpha_N^{(1)}(0) = \frac{Np(1-p)^{N-1}}{1-(1-p)^N}$$

Taking the derivative of $\alpha_N^{(1)}(0)$ wrt N gives

$$\frac{p(1-p)^{N-1} [1 - (1-p)^N + N \ln(1-p)]}{[1 - (1-p)^N]^2}. \quad (10)$$

The sign of (10) depends on $1 - (1-p)^N + N \ln(1-p)$, which decreases in p . Setting $p = 0$ in this expression gives 0, which implies that (10) is always negative. As a result $\alpha_N^{(1)}(0)$ decreases in N (first part of Assumption 1). Moreover, the mean of the TBD is $Np/[1 - (1-p)^N]$. Taking the derivative of the mean wrt N gives

$$\frac{p - p(1-p)^N [1 - N \ln(1-p)]}{[1 - (1-p)^N]^2}. \quad (11)$$

The sign of (11) depends on $1 - (1-p)^N [1 - N \ln(1-p)]$, which increases in N . Setting $N = 2$ in this expression gives $1 - (1-p)^2 [1 - 2 \ln(1-p)] > 0$ for all p . Hence (11) is always positive and therefore the mean increases in N (second part of Assumption 1).

We now consider condition (8). We have that

$$\frac{\alpha_N^{(1)}(0)}{\alpha_N^{(1)}(1-\tau)} = \frac{\frac{Np(1-p)^{N-1}}{1-(1-p)^N}}{\frac{Np(1-p\tau)^{N-1}}{1-(1-p\tau)^N}} = \left(\frac{1-p}{1-p\tau} \right)^{N-1},$$

which clearly decreases in N . As a result, when the distribution of price information in the market follows the TBD, an increase in the number of competitors leads to lower (in a FOSD sense) prices.

In order to understand whether lower percentiles fall more in N than higher percentiles, we take the derivative of (7)

$$\frac{\alpha_{N+1}^{(1)}(0)}{\alpha_{N+1}^{(1)}(1-\tau)} - \frac{\alpha_N^{(1)}(0)}{\alpha_N^{(1)}(1-\tau)} = \left(\frac{1-p}{1-p\tau} \right)^N - \left(\frac{1-p}{1-p\tau} \right)^{N-1}$$

wrt τ , which gives

$$\left(\frac{1-p}{1-p\tau} \right)^{N-2} \frac{p(1-p)}{(1-p\tau)^3} [1 - Np(1-\tau) - p\tau]. \quad (12)$$

Inspection of (12) reveals that its sign is always positive provided that $\tau > (Np - 1)/[p(N - 1)]$. From this we conclude that when $p < 1/N$, (12) is positive for all τ , hence the lower percentiles fall more than higher percentiles. When $p > 1/N$, then there exists a critical $\check{\tau}$ such that the fall in the percentiles increases in $[0, \check{\tau}]$ and decreases in $[\check{\tau}, 1]$. Figure 1 presents two examples.

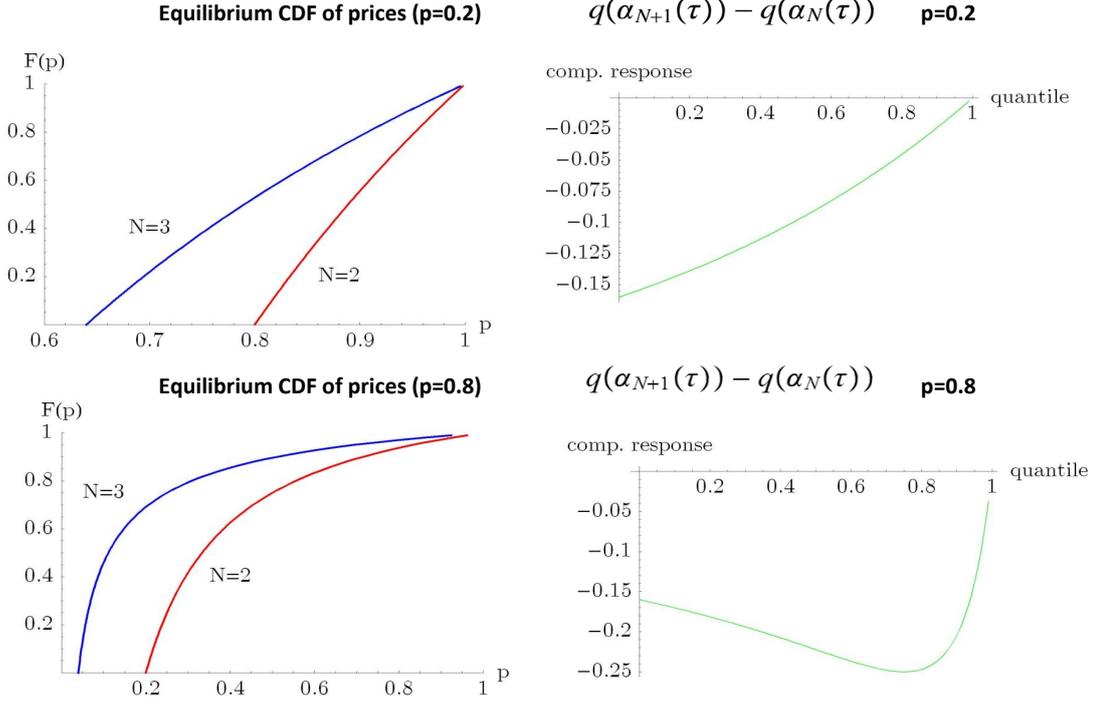


Figure 4: The truncated binomial distribution

Example 2 (Varian's (1980) information structure) *Varian's (1980) information structure has PGF*

$$\alpha_N(x) = \mu x + (1 - \mu)x^N,$$

for some $0 < \mu < 1$. Note that $\alpha_N^{(1)}(0) = \mu$, which is constant in N , and the mean is $\mu + N(1 - \mu)$, which increases in N . Therefore, Assumption 1 holds.

Regarding condition (8), we have

$$\frac{\alpha_N^{(1)}(0)}{\alpha_N^{(1)}(1 - \tau)} = \frac{\mu}{\mu + N(1 - \mu)(1 - \tau)^{N-1}} \quad (13)$$

Taking the derivative of (13) wrt N gives

$$-\frac{\mu(1 - \mu)(1 - \tau)^{N-1} [1 + N \ln(1 - \tau)]}{[\mu + N(1 - \mu)(1 - \tau)^{N-1}]^2}. \quad (14)$$

The sign of this expression is the opposite of the sign of $1 + N \ln(1 - \tau)$, which is positive for $\tau \leq \exp[-1/N]$ and negative otherwise. As a result, (14) is negative for low τ and positive for high τ . We conclude that condition (8) is violated.

Consider now condition (9). Since $\alpha_{N+1}^{(s)}(0) = 0$ for all $s = 2, \dots, N$, we invoke the $(N + 1)^{th}$

derivative of $\alpha_{N+1}(x)$. Then we have

$$\frac{\alpha_{N+1}^{(N+1)}(0)}{\alpha_{N+1}^{(1)}(0)} = \frac{(N+1)!(1-\mu)}{\mu} > 0,$$

which implies that condition (9) holds. We conclude that when the distribution of price information in the market follows Varian's distribution, an increase in the number of competitors leads to an increase in the high percentiles of the price distribution, and to a decrease in the low percentiles of the price distribution. The top panels in Figure 3 provide an example (for $\mu = 0.5$).

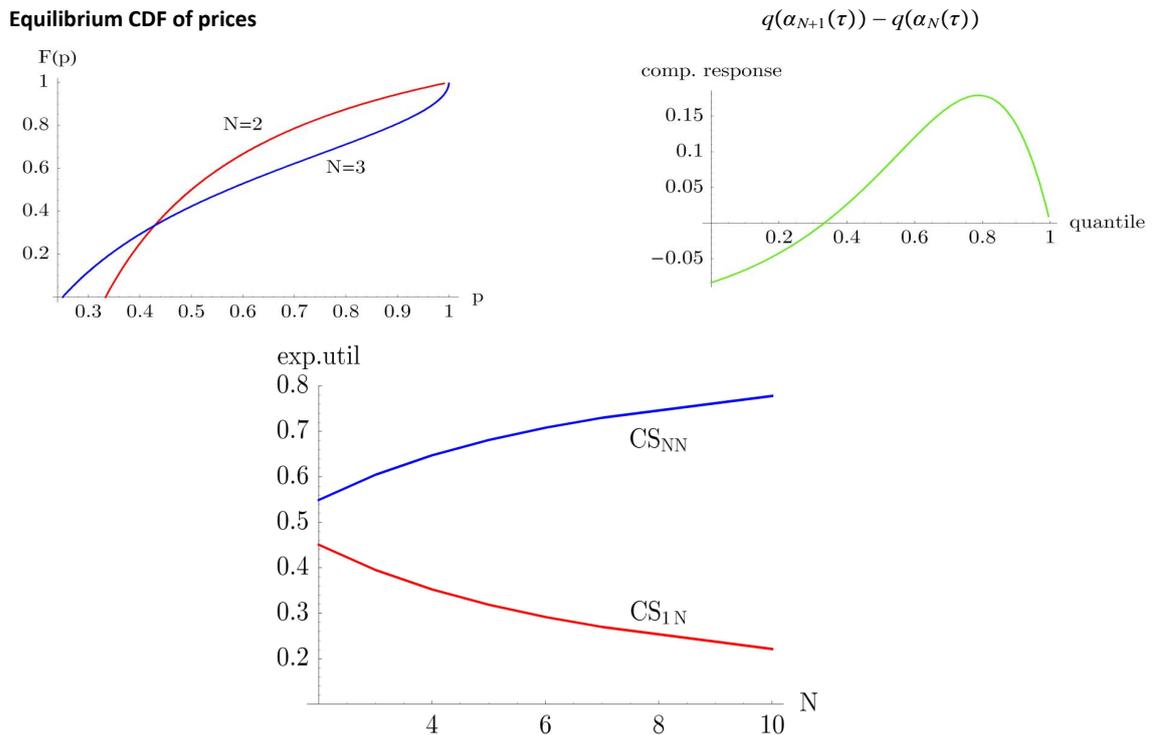


Figure 5: The shopper/non-shopper distribution ($\mu = 0.5$)

3.2 Consumer welfare and the number of firms

We have seen that the response to an increase in competition differs across the percentiles of the price distribution. In particular, we have shown (i) that some percentiles may increase while others decrease, and (ii) that if they all fall, some may decrease more than others, depending on the properties of the PGF $\alpha_N(x)$. These two results have important implications. Suppose, first, that all percentiles decrease when the number of competitors increases from N to $N+1$. Because some prices may decline more than others, it is possible that consumers observing a given number of prices derive greater benefits from increased competition than other consumers observing a different number of

prices. Second, because the frequency of high prices can actually increase, some consumers may end up paying higher prices, even on average, after an increase in the number of competitors. These implications of the model are in stark contrast to standard full-information oligopoly models. In this subsection we proceed to study the welfare gains that different consumers (i.e., consumers observing different number of prices) will derive from an increase in competition.⁴³

The utility of a consumer who buys from a firm i at a price p_i is given by $v - p_i$. Denote the utility of a consumer who observes s prices and buys from the cheapest seller by $u_s = \max \{v - p_1, v - p_2, \dots, v - p_s\}$ where p_1, p_2, \dots, p_s are i.i.d. random variables drawn from the equilibrium price distribution F . The distribution of u_s is $(1 - F(v - u))^s$. Using the price distribution from Section 2, we can derive the inverse of the distribution of u_s :

$$y_s(\alpha_N(\tau)) = k \left[1 - \frac{\alpha_N^{(1)}(0)}{\alpha_N^{(1)}(\tau^{1/s})} \right]. \quad (15)$$

where $\tau \in [0, 1]$. In this case, for a given τ , (15) gives the τ^{th} percentile of the distribution of the maximum utility received by a consumer who observes s prices. Because v is fixed this distribution provides the same information as the distribution of prices paid and we will also informally refer to (15) as the τ^{th} percentile of the distribution of prices paid.

Following the same steps as before, we can study how the distribution of utilities received by a consumer who observes s prices changes when we move from an N -firm to an $N + 1$ -firm market. We have

$$y_s(\alpha_{N+1}(\tau)) - y_s(\alpha_N(\tau)) = k \left[\frac{\alpha_N^{(1)}(0)}{\alpha_N^{(1)}(\tau^{1/s})} - \frac{\alpha_{N+1}^{(1)}(0)}{\alpha_{N+1}^{(1)}(\tau^{1/s})} \right] \quad (16)$$

Note the similarity between the expression in (16) and that in equation (7). In our empirical section we take advantage of this equation in order to estimate the impact of an increase in N on the utility distribution of a consumer observing s prices. Comparing (16) and (7), it is immediate to see that the change in the percentile τ of the price distribution is equivalent to the change in the percentile $(1 - \tau)^s$ of the utility distribution of a consumer observing s prices. We then have:

Corollary 2 (of Proposition 2) *Suppose that the number of firms increases from N to $N + 1$ and that Assumption 1 holds. Then, for all $s = 1, 2, \dots, N$:*

(I) *There exists a percentile $\hat{\tau}_s \in [0, 1]$ such that all the percentiles $\tau > \hat{\tau}_s$ of the distribution of utilities received by a consumer observing s prices increase.*

⁴³Note that we ignore non-price effects of competition such as better quality of service, shorter distances to retailers, etc., so that “welfare effects” here refers exclusively to price effects.

(II) All the percentiles of the distribution of utilities received by a consumer observing s prices increase if and only if condition (8) holds.

(III) If condition (9) holds, then there exists $\tilde{\tau}_s \in (0, 1]$ such that all the percentiles below $\tilde{\tau}_s$ of the distribution of utilities received by a consumer observing s prices decrease.

In line with Proposition 2, an increase in the number of firms results in an increase in the upper percentiles of the distribution of utilities derived by all consumers. When condition (8) holds, then all the percentiles of the utility distribution of a given consumer group increase after the number of firms goes up. Under condition (9), the distribution of the price paid by any type of consumer when there are $N + 1$ firms in the market surely “crosses-over” the distribution function when there are N firms. In this situation, increased competition may, like in Varian’s (1980) model, raise the price paid by consumers who pay a high price, despite Assumption 1 being satisfied.⁴⁴

This result states that, under condition (8), all consumers will obtain greater expected utilities given the number of prices observed. The proposition however does not inform us about whether some consumers benefit more than others. In fact, it is quite difficult to evaluate analytically how the change in consumer surplus

$$CS_{sN+1} - CS_{sN} = k \int_0^1 \left[\frac{\alpha_N^{(1)}(0)}{\alpha_N^{(1)}(\tau^{1/s})} - \frac{\alpha_{N+1}^{(1)}(0)}{\alpha_{N+1}^{(1)}(\tau^{1/s})} \right] d\tau \quad (17)$$

depends on s . Using the examples above, however, we can gain some insights into this issue. For example, for the TBD case with $p < 1/N$, the lower percentiles fall more than the higher ones so we expect utility gains from increased competition to rise in s . By contrast, when $p > 1/N$, the impact of increased competition is felt more at intermediate percentiles of the equilibrium price distribution and as a result we expect utility gains from increased competition to fall in s .

Example 1 (cont’d) (The truncated binomial distribution). *When the price information consumers have follows the TBD, the utility gains from increased competition, equation (17), are as follows*

	<i>Expected Utility Gains ($p = 0.2$)</i>		
	$N = 2$	$N = 3$	$N = 4$
$s = 1$	0.092	0.080	0.069
$s = 2$	0.118	0.101	0.085
$s = 3$		0.109	0.092
$s = 4$			0.095

	<i>Expected Utility Gains ($p = 0.8$)</i>		
	$N = 2$	$N = 3$	$N = 4$
$s = 1$	0.202	0.080	0.037
$s = 2$	0.198	0.061	0.021
$s = 3$		0.051	0.015
$s = 4$			0.012

⁴⁴Anecdotal evidence tells us that many consumers report not to have felt the “supposed gains” from increased competition in liberalized markets such as airlines, gasoline, telecoms, etc. This might be related to this fact in combination with consumers remembering bad news (about prices) more readily than good news.

In this case, gains from increased competition increase in s when p is small and decrease in s when p is high.

This example shows two important points: (i) consumer benefits from increased competition depend on s and therefore utility gains are asymmetrically distributed across the consumer population; and (ii) utility gains need not be increasing in s , that is, poorly informed consumers may benefit more than well informed ones.

Corollary 2 also states the remarkable result that some consumers may experience a welfare loss after the number of competitors increases. In fact, when condition (9) holds, $\alpha_{N+1}^{(1)}(0) = \alpha_N^{(1)}(0)$ suffices for those consumers observing one price only to experience a welfare loss. This explains why mean prices increase in N in Varian (1980).

Example 2 (cont'd) (Varian's (1980) distribution of information). *When the distribution of price information among consumers follows Varian's distribution then the utility gains from increased competition are negative for the consumers who observe one price only and positive for the consumers who observe all prices in the market. The bottom panel in Figure 3 provides an example for $\mu = 0.5$.*

4 Conclusions

In markets where the amount of price information consumers have is heterogeneous, prices are typically dispersed in equilibrium. When this happens, the relevant question is no longer what happens to the price when the number of firms changes but, instead, what happens to the whole distribution of equilibrium prices. Using data from the gasoline market in the Netherlands, we have found, first, that markets with a given number of competitors have price distributions that first-order stochastically dominate the corresponding price distributions in markets with one more firm. Second, our data have revealed that the competitive response varies along the price distribution and is stronger at prices in the medium to upper part of the distribution. Finally, we have performed simulations to uncover how the gains from increased competition vary as consumer price information changes. It turns out that consumers who are less alert to the posted prices by the firms derive lower gains from competition than those who are more attentive and stay well informed.

To account for these empirical results, we have proposed a generalisation of Varian's (1980) well-known model of sales that allows for richer heterogeneity in consumer price information. The model makes clear that increased competition has an effect on prices only when it changes the amount of price information consumers have. Though the generalisation increases the complexity of the model,

we have shown that it can generate the observed patterns in the data. We view this as an important observation because Varian's model distinct predictions have sometimes been used to dismiss it as a possible explanation for observed price dispersion in real-world markets.

Since price dispersion is prevalent in many markets, we believe the paper has a general message that goes beyond the present application to the gasoline market in the Netherlands. The price effects of competition-enhancing policies (e.g., industry deregulation, trade liberalization, merger control, etc.) are not as straightforward as one may be led to believe based on standard oligopoly theory. As a result, welfare implications are not obvious either. In fact, we have shown, theoretically and empirically, that increased competition can have unequal effects among consumers; at least theoretically, some consumers may even experience declines in their welfare as a result of some prices going up.

In order to identify which consumers benefit more and which benefit less from increased competition in gasoline prices we would require a mapping between shopping behavior and socio-economic characteristics of interest. If such data were available we would be able to say something about how the distribution of the benefits from increased competition varies with income. For example, consumers that observe only a few prices may be high-income consumers (whose value of time is higher) and these consumers may benefit less from competition than low-income consumers. This, however, is beyond the scope of this paper and is left for future research. Moreover, a complete welfare analysis, should recognize the effect of increased competition on other dimensions of consumer welfare such as increased variety, quality and accessibility.

Lastly, we think the empirical findings reported in the paper and that have led us to extend Varian's model are of interest on their own right and, if verified in other data sets, they should be taken into account when formulating theoretical models of pricing in oligopolistic markets.

References

- [1] Armstrong, Mark (2008), “Interactions between Competition and Consumer Policy,” *Competition Policy International*, 4(1), 97-147.
- [2] Armstrong, Mark, John Vickers and Jidong Zhou (2009), “Consumer Protection and the Incentive to Become Informed”, *Journal of the European Economic Association* (Papers and Proceedings), 7(2-3), 399-410.
- [3] Barut, Yasar and Dan Kovenock (1998), “The Symmetric Multiple Prize All-Pay Auction with Complete Information,” *European Journal of Political Economy*, 14, 627-644.
- [4] Baye, Michael. R., Dan Kovenock, and Casper de Vries (1992), “It Takes Two to Tango: Equilibria in a Model of Sales,” *Games and Economic Behavior*, 4, 493–510.
- [5] Baye, Michael. R., John Morgan and Patrick Scholten (2004), “Price Dispersion in the Small and in the Large: Evidence from an Internet Price Comparison Site”, *The Journal of Industrial Economics*, 52(4), 463-496.
- [6] Baye, Michael. R., John Morgan and Patrick Scholten (2006), “Information, Search, and Price Dispersion, chapter 6 in *Handbook on Economics and Information Systems Volume 1*, (T. Hendershott, Ed.), Amsterdam: Elsevier.
- [7] Barron, John M., Beck A. Taylor and John R. Umbeck (2004), “Number of sellers, average prices, and price dispersion”, *International Journal of Industrial Organization*, 22(8-9), 1041-1066.
- [8] Borenstein, Severin and Nancy Rose (1994), “Competition and Price Dispersion in the U.S. Airline Industry,” *Journal of Political Economy*, 102, 653-683.
- [9] Burdett, Kenneth and Kenneth L. Judd (1983), “Equilibrium Price Dispersion,” *Econometrica*, 51, 955-70.
- [10] Chandra, Ambarish and Mariano Tapatta (2011), “Consumer Search and Dynamic Price Dispersion: An Application to Gasoline Markets,” *The Rand Journal of Economics*, 42(4), 681–704.
- [11] Chernozhukov, Victor (2005), “Extremal Quantile Regression,” *Annals of Statistics*, 33(2), 806–839.

- [12] Dana, James D. (1994), “Learning in an Equilibrium Search Model,” *International Economic Review*, 35(3), 745–771.
- [13] Janssen, Maarten and José L. Moraga-González (2004), “Strategic Pricing, Consumer Search and the Number of Firms”, *Review of Economic Studies*, 71(4), 1089–1118.
- [14] Galeotti, Andrea (2010), “Searching, Talking and Pricing”, *International Economic Review*, 51(4), 1159–1174.
- [15] Gerardi, Kristopher S. and Adam Hale Shapiro (2009), “Does Competition Reduce Price Dispersion: New Evidence from the Airline Industry”, *Journal of Political Economy*, 117(1), 1–37.
- [16] Hagemann, Andreas (2015), “Cluster-robust bootstrap inference in quantile regression models”, *Journal of the American Statistical Association*, forthcoming.
- [17] Hastings, Justine S. (2004), “Vertical Relationships and Competition in Retail Gasoline Markets: Empirical Evidence from Contract Changes in Southern California,” *American Economic Review*, 94, 317–328.
- [18] Haynes, Michelle and Steve Thompson (2008), “Price, Price Dispersion and Number of Sellers at a Low Entry Cost Shopbot,” *International Journal of Industrial Organisation*, 26, 459-72.
- [19] Hortagsu, Ali and Chad Syverson (2004), “Product Differentiation, Search Costs and Competition in the Mutual Fund Industry: A Case Study of S&P 500 Index Funds”, *Quarterly Journal of Economics*, 119, 403–456.
- [20] Hosken, Daniel S., Robert S. McMillan and Christopher T. Taylor (2008), “Retail Gasoline Pricing: What Do We Know?”, *International Journal of Industrial Organization*, 26(6), 1425-1436.
- [21] Houde, Jean-François (2012), “Spatial Differentiation and Vertical Mergers in Retail Markets for Gasoline,” *American Economic Review*, 102(5), 2147–2182.
- [22] Koenker, Roger (2004), “Quantile regression for longitudinal data,” *Journal of Multivariate Analysis*, 91(1), 74-89.
- [23] Lach, Saul (2002), “Existence and Persistence of Price Dispersion: an Empirical Analysis”, *Review of Economics and Statistics*, 84(3), 433-444.

- [24] Lewis, Matthew (2008), “Price Dispersion and Competition with Differentiated Sellers”, *Journal of Industrial Economics*, 56 (3), 654–678.
- [25] Marvel, H.P. (1976), “The Economics of Information and Retail Gasoline Price Behavior: An Empirical Analysis,” *Journal of Political Economy*, 84, 1033–1060.
- [26] Milgrom, Paul R. (1981), “Good News and Bad News: Representation Theorems and Applications,” *The Bell Journal of Economics*, 12(2), 380–391.
- [27] Moldovanu, Benny and Aner Sela (2001), “The Optimal Allocation of Prizes in Contests,” *American Economic Review*, 91(3), 542–558.
- [28] Moraga-González, José L., Zsolt Sándor and Matthijs R. Wildenbeest (2016), “Prices and Heterogeneous Search Costs”, *The Rand Journal of Economics*, forthcoming.
- [29] Morgan, John, Henrik Orzen and Martin Sefton (2006), “An experimental study of price dispersion,” *Games and Economic Behavior*, 54, 134–158.
- [30] Nermuth, Manfred, Giacomo Pasini, Paolo Pin and Simon Weidenholzer (2009), “The Informational Divide”, *Games and Economic Behavior*, 78, 2013, 21-30.
- [31] NMa (2006): Benzinescan 2005/2006, De Nederlandse Mededingingsautoriteit.
- [32] Pennerstorfer, Dieter, Philipp Schmidt-Dengler, Nicolas Schutz, Christoph Weiss and Biliana Yontcheva (2015), “Information and Price Dispersion: Evidence from Retail Gasoline,” unpublished manuscript.
- [33] Rosenthal, Robert W. (1980), “A Model in Which an Increase in the Number of Sellers Leads to a Higher Price”, *Econometrica*, 48, 1575–1579.
- [34] Shepard, A. (1991), “Price discrimination and retail configuration,” *Journal of Political Economy*, 99, 30–53.
- [35] Soetevent, Adriaan R., Marco A. Haan, and Pim Heijnen (2014), “Do Auctions and Forced Divestitures Increase Competition?,” *The Journal of Industrial Economics*, 62(3), 467–502.
- [36] Varian, Hal (1980), “A Model of Sales,” *American Economic Review*, 70, 651–659.
- [37] Waterson, Michael (2003), “The Role of Consumers in Competition and Competition Policy,” *International Journal of Industrial Organization*, 21, 129–50.

- [38] Wildenbeest, Matthijs R. (2011), An empirical model of search with vertically differentiated products, *The Rand Journal of Economics*, 42, 729-57.
- [39] Wu, Ximing and Jeffrey M. Perloff (2007), "Tax Incidence Varies Across the Price Distribution," *Economics Letters* 96, 93-99.

Appendix: Proofs

Proof of Proposition 2. First we note that

$$q(\alpha_{N+1}(\tau)) - q(\alpha_N(\tau)) = k \left[\frac{\alpha_{N+1}^{(1)}(0)}{\alpha_{N+1}^{(1)}(1-\tau)} - \frac{\alpha_N^{(1)}(0)}{\alpha_N^{(1)}(1-\tau)} \right] = \frac{k}{h(\tau)} \left[\frac{\alpha_N^{(1)}(1-\tau)}{\alpha_N^{(1)}(0)} - \frac{\alpha_{N+1}^{(1)}(1-\tau)}{\alpha_{N+1}^{(1)}(0)} \right] \quad (18)$$

where

$$h(\tau) \equiv \frac{\alpha_{N+1}^{(1)}(0)\alpha_N^{(1)}(0)}{\alpha_{N+1}^{(1)}(1-\tau)\alpha_N^{(1)}(1-\tau)} > 0.$$

(I) Let us set $\tau = 0$ in this expression. It follows that the sign of $q(\alpha_{N+1}(0)) - q(\alpha_N(0))$ equals the sign of

$$-\frac{1}{\alpha_N^{(1)}(1)\alpha_{N+1}^{(1)}(1)} \left[\alpha_{N+1}^{(1)}(1)\alpha_N^{(1)}(0) - \alpha_N^{(1)}(1)\alpha_{N+1}^{(1)}(0) \right] < -\frac{1}{\alpha_{N+1}^{(1)}(1)} \left[\alpha_N^{(1)}(0) - \alpha_{N+1}^{(1)}(0) \right] \leq 0,$$

where the two inequalities follow from Assumption 1. Since (18) is strictly negative when $\tau = 0$, by continuity of the function $q(\alpha_{N+1}(\tau)) - q(\alpha_N(\tau))$ in τ , we conclude that the low percentiles of the price distribution will always decrease.

(II) If condition (8) holds, then it is straightforward to see that all percentiles will fall.

(III) To prove this, we first note that setting $\tau = 1$ in (18) gives 0. Now, let us take the derivative of (18) wrt τ . We get:

$$\frac{\partial}{\partial \tau} (q(\alpha_{N+1}(\tau)) - q(\alpha_N(\tau))) = k \left\{ \frac{\alpha_{N+1}^{(1)}(0)\alpha_{N+1}^{(2)}(1-\tau)}{[\alpha_{N+1}^{(1)}(1-\tau)]^2} - \frac{\alpha_N^{(1)}(0)\alpha_N^{(2)}(1-\tau)}{[\alpha_N^{(1)}(1-\tau)]^2} \right\}$$

Setting $\tau = 1$ gives

$$\frac{\partial}{\partial \tau} (q(\alpha_{N+1}(0)) - q(\alpha_N(0))) = k \left[\frac{\alpha_{N+1}^{(2)}(0)}{\alpha_{N+1}^{(1)}(0)} - \frac{\alpha_N^{(2)}(0)}{\alpha_N^{(1)}(0)} \right] \quad (19)$$

When condition (9) holds, this expression is negative. This implies that the difference $q(\alpha_{N+1}(\tau)) - q(\alpha_N(\tau))$ is decreasing in a neighborhood of $\tau = 1$. Since it is zero when $\tau = 1$, by continuity we conclude that $q(\alpha_{N+1}(\tau)) > q(\alpha_N(\tau))$ for sufficiently large percentiles.⁴⁵ ■

⁴⁵It could be the case that (19) is zero. In that case, it is straightforward to see that higher order derivatives can be invoked. For example, we can take the second order derivative of (18) and evaluate it at $\tau = 1$, which gives

$$k \left[\frac{\alpha_{N+1}^{(3)}(0)}{\alpha_{N+1}^{(1)}(0)} - \frac{\alpha_N^{(3)}(0)}{\alpha_N^{(1)}(0)} \right].$$

If this expression is negative, then the derivative of (18) decreases in a neighborhood of $\tau = 1$. As a result $q(\alpha_{N+1}(\tau)) - q(\alpha_N(\tau))$ must be concave in a neighborhood of $\tau = 1$ and since it itself and its derivative are equal to zero at $\tau = 1$, we conclude $q(\alpha_{N+1}(\tau)) - q(\alpha_N(\tau)) > 0$ in a neighborhood of $\tau = 1$.

Proof of Corollary 1. Using (18), we have that $q(\alpha_{N+1}(\tau)) - q(\alpha_N(\tau)) < 0$ if and only

$$\begin{aligned} \frac{\alpha_N^{(1)}(1-\tau)}{\alpha_N^{(1)}(0)} - \frac{\alpha_{N+1}^{(1)}(1-\tau)}{\alpha_{N+1}^{(1)}(0)} &= \sum_{s=1}^N s \frac{\mu_s(N)}{\mu_1(N)} (1-\tau)^{s-1} - \sum_{s=1}^{N+1} s \frac{\mu_s(N+1)}{\mu_1(N+1)} (1-\tau)^{s-1} \\ &= \sum_{s=1}^N s \left[\frac{\mu_s(N)}{\mu_1(N)} - \frac{\mu_s(N+1)}{\mu_1(N+1)} \right] (1-\tau)^{s-1} - (N+1) \frac{\mu_{N+1}(N+1)}{\mu_1(N+1)} (1-\tau)^N < 0. \end{aligned}$$

From this expression, it is clear that

$$\frac{\alpha_N^{(s)}(0)}{\alpha_N^{(1)}(0)} - \frac{\alpha_{N+1}^{(s)}(0)}{\alpha_{N+1}^{(1)}(0)} = s! \left[\frac{\mu_s(N)}{\mu_1(N)} - \frac{\mu_s(N+1)}{\mu_1(N+1)} \right] \leq 0 \text{ for all } s \text{ suffices.}$$

■

Proof of Corollary 2. Follows straightforwardly from Proposition 2. ■