

Product differentiation decisions under ambiguous consumer demand and pessimistic expectations*

Michał Król[†]

ABSTRACT. This paper studies product differentiation decisions in a spatial duopoly with limited information on consumer demand. In particular, a situation is discussed in which the firms do not know the exact distribution of the random location of consumer demand and its responsiveness to price changes (measured by the scale of transport costs), but resolve the resulting ambiguity using the α -maxmin or minmax regret criteria. When the firms are sufficiently pessimistic (α is high enough), results are in contrast with the existing literature. In particular, an increase of demand location uncertainty decreases the equilibrium product differentiation, intensifying the second-stage competition in prices, although the effect is dampened by uncertainty about transport costs. Endogenizing the choice of objective function leads to the dominance of an extreme form of pessimism, which turns out to be socially-optimal.

Keywords: Hotelling, uncertainty, non-unique prior, pessimism

JEL Classification: C72, D43, D81, L13, R32

1 Introduction

Over the last few decades, the Hotelling model of spatial competition has been used to explain a wide variety of social phenomena, from the location of retail outlets to competition among political parties. A relatively new strand of the relevant literature investigates the effects of random demand fluctuations on the firms' location decisions. This typically entails introducing some form of demand uncertainty into a modified Hotelling setting. For

*I would like to thank Prof. Paul Madden for all his support, and Prof. Roger Hartley and Prof. Leo Kaas for their comments on the paper, as well as two anonymous reviewers and the editor, Prof. Bruno Jullien, for their helpful suggestions. All the remaining errors are my own. This work was supported by the Economic and Social Research Council [grant number ES/G016321/1].

[†]Department of Economics, The University of Manchester, michal.krol@manchester.ac.uk

instance, Balvers and Szerb [1996] consider the effect of random shocks to the products' quality / desirability when prices are fixed. The last restriction is relaxed by Harter [1997], who studies uncertainty as a uniformly distributed random shift of the (uniform) customer distribution, given that firms locate sequentially.

Relatively few studies are concerned with the effect of demand uncertainty in an otherwise unchanged Hotelling framework. In particular, Casado-Izaga [2000] uses the same form of uncertainty as Harter, but the duopolists locate simultaneously, before observing the customer distribution and only then naming prices. This setting is generalized by Meagher and Zauner [2005], who parametrize the support of the (uniform) random variable that shifts the customer distribution. They report that demand uncertainty increases the equilibrium level of product differentiation. In another study (henceforth, MZ), Meagher and Zauner [2004] consider a random shock arbitrarily, rather than uniformly, distributed on a fixed interval. Tractability of the model is maintained by assuming that the variance of the shock is small enough relative to the ex-post differentiation of tastes, so that no firm would ever choose to capture the entire market in any state of nature. Once again, it turns out that more uncertainty results in a higher equilibrium level of product differentiation.

The intuition for those results is simple. If the demand is more likely to be located away from the centre of the market, then it is natural for the firms to venture into more distant areas and away from one another, relaxing the second-stage price competition.

Nevertheless, all the studies above rely on the common prior assumption. Not only are the firms able to estimate the probability distribution of random factors affecting their profits, but they also arrive in this respect at exactly the same conclusions. This may often be impossible in reality, particularly at times when the scale of market disturbances is indicative of structural changes in the patterns of consumer choice, rather than merely random fluctuations around an otherwise unchanged standard of behaviour. In such crises, as well as at the opening of new markets, sellers are confronted with an ambiguously distributed demand, until they are able to gather sufficient information, e.g. via market research practices or price experimentation (see Aghion et al. [1993]).

Faced with such ambiguity, firms might at least be able to place the potential demand variations within certain bounds. For instance, in the classic 'Main Street' example, a sudden influx of tourists to the city could make retailers ignorant of the likelihood that consumers will cluster in one part of the street or the other. Even so, sellers might still reasonably expect their customers to be confined to the physical boundaries of the Main Street. Similarly, many product characteristics, such as sugar content in food, are naturally restricted by certain bounds (in this case zero and one hundred percent).

The purpose of this paper is then to investigate how the firms' product differentiation decisions would change if the information on potential demand fluctuations was limited in the above manner. In addition, since the absence of a unique prior precludes expected profit maximization, it would be interesting to see how the choice of objective function used as an alternative affects the outcome of the firms' interaction.

In particular, for a given location decision the duopolists are able to identify the worst-

and best-case scenarios within the range of possible demand variations. Hence, one natural objective is to maximize a weighted average of the associated extreme profit values, i.e. use the Arrow and Hurwicz [1972] alpha-maximin criterion to resolve the ambiguity, where the weight assigned to the lowest possible profit can be said to constitute a measure of the player’s pessimism. It turns out that, contrary to the existing literature, firms that are sufficiently pessimistic in this sense locate closer together in response to an increase of uncertainty (represented by a spread of the range of possible demand variations). This is because an increase in the scale of potential demand fluctuations means a player has to consider a possibility of the worst-case demand shifting even more to her disadvantage (and in favour of the competitor). To reduce one’s exposure to this threat, a pessimistic player will then make her product resemble that of the rival, leaving less room for becoming handicapped, even at the cost of less product differentiation inducing more competitive pricing.

As detailed in section 5, this could explain some of the observable industrial behaviour. For instance, due to their traditional objective of ‘balancing the odds’, bookmakers effectively concentrate on the lowest-profit scenario, in which the outcome of a sporting event that attracted the largest volume of bets is realized. Similarly, many financial institutions focus on the worst-case scenario so as to satisfy the self-imposed or government enforced stress tests. In both of these situations, evidence can be given of reduced product differentiation as a result of an increase in uncertainty.

In contrast, it is shown that applying an alternative minmax regret objective function, as initially proposed by Savage [1951], reproduces the existing ‘common prior’ comparative statics results. This is in line with the work of Bergemann and Schlag [2008a,b], who studied monopolistic pricing by sellers aware of only the support (but not the exact distribution) of buyers’ valuations. They find that more uncertainty reduces prices given a maximin objective, but not necessarily under the minmax regret scheme.

To resolve the present discrepancy between the alternative payoff specifications, in Section 4 I consider a possibility of the firms choosing their objective function (out of the alpha-maximin class, as well as the minmax regret), before moving on to location decisions. It turns out that focusing exclusively on the worst-case scenario is a dominant strategy *regardless of the resolution of the uncertainty*. This is in contrast with existing studies on optimism in trading/investment (e.g. Kyle and Wang [1997], Englmaier [2010]), as well as in more general competitive contexts (Heifetz et al. [2007a,b]), which predict optimism becoming prevalent due to the pessimistic agents being outperformed.

Section 6 discusses and further elaborates on the above results. Firstly, out of all the considered objective functions, the outcome associated with the most pessimistic (maximin) approach is shown to be not only the most competitive, but also the most socially-desirable one, whatever the actual distribution of the random shock. Furthermore, it is demonstrated that the same competitive outcome may be approximated under probabilistic information via an optimally set progressive corporate taxation scheme, making the firms effectively risk-averse. What is then interesting about ambiguity is that it makes such a government

intervention unnecessary due to the pessimistic attitudes becoming prevalent. However, a distinguishing factor between risk aversion under probabilistic information and pessimism under ambiguity is how uncertainty about the demand location interacts with that about other aspects of the demand, such as its sensitivity to price changes. This is represented by the scale of transportation costs¹, which has no effect on location decisions under certainty. Nevertheless, uncertainty about costs makes the outcome less competitive under ambiguity, while with probabilistic information it can have an opposite effect. Finally, Section 6 concludes with a discussion of robustness of the obtained results to changes in the timing of moves or the transport cost specification.

2 The Model

The setting is, in general, very similar to Meagher and Zauner [2004] (MZ). In the first stage of the game, two firms simultaneously choose locations x_1, x_2 (without loss of generality set $x_1 \leq x_2$) and then proceed to simultaneous setting of their respective prices p_1, p_2 in the second stage². As usual, a consumer located at x chooses to buy a unit of the good from firm $i \in \{1, 2\}$, so as to minimize the total purchase cost of $p_i + t(x_i - x)^2$, where $t > 0$ is the transportation cost parameter. The good costs nothing to produce and the consumers are uniformly distributed on the interval $[M - \frac{1}{2}, M + \frac{1}{2}]$, where the duopolists get to know the value of M , as well as t , once they choose the locations, but before setting prices. Initially, all they know is that the joint probability distribution of (M, t) has support $S \equiv [-L, L] \times [t_0, 1]$, where $L \in [0, \frac{1}{2}]$ ³ and $t_0 \in (0, 1]$. Note that the upper bound of t is set to 1 without loss of generality, while the assumption that $t > 0$ is present in all product differentiation models (otherwise, the consumers only consider prices and location / product design decisions are inconsequential). It should also be emphasized that, as in MZ, the two firms have identical information, which in this case means they agree on what the support S is. In reality, it may at times be the case that firms have different beliefs in this respect, but this possibility is not pursued here.

The first difference from MZ is introducing uncertainty about transportation costs. MZ assumes $t = 1$, which may be done without loss of generality, so long as the scale of transportation costs is constant. However, with probabilistic information uncertainty about costs may interact with that about the demand location and influence location decisions,

¹The total consumer demand is, by assumption, completely inelastic in the Hotelling framework. However, when the transport cost parameter decreases, the individual demand of each firm for given locations and the counterpart's price becomes more elastic in the firm's own price.

²Note that the analysis is restricted to pure strategies in both locations and prices - see Bester et al. [1996] for a study of mixed strategy location choice, and Osborne and Pitchik [1987] for an analysis of mixed strategy pricing.

³The assumption that $L \leq 1/2$ was imposed in MZ for the purpose of mathematical tractability and is equally useful here.

just as it turns out to be the case in the present model. See Section 3 and Section 6 for a more detailed discussion of this issue.

The second difference from MZ is the fact that the exact probability distribution of (M, t) is unknown, and so is the expected value of the second stage pure-strategy Nash Equilibrium profits (which continue to be uniquely determined for any location-pair and any outcome of the uncertainty). Hence, in their location decisions, the duopolists must pursue an objective different from expected profit maximization.

One possibility is to consider a weighted average of the lowest and highest possible profits, i.e. to use the Arrow/Hurwicz α -maxmin criterion instead of the expected value. More specifically, let $\pi_i^*(x_i, x_{-i}, M, t)$ be the unique second-stage Nash Equilibrium profit associated with a particular location-pair and demand realization. Then the payoffs in the reduced first-stage game are given by:

$$\Pi_i^\alpha(x_1, x_2) = \alpha \left[\min_{(M,t) \in S} \pi_i^*(x_i, x_{-i}, M, t) \right] + (1 - \alpha) \left[\max_{(M,t) \in S} \pi_i^*(x_i, x_{-i}, M, t) \right]$$

where $\alpha \in [0, 1]$ is a parameter representing the degree of a duopolist's pessimism.

In fact, the above specification coincides with the α -Maxmin Expected Utility (see Ghirardato et al. [2004]) of a risk-neutral agent. To see this, suppose the agent's utility function is simply $u(x) = x$, and the set Φ of priors includes all possible probability measures on S , i.e. $\Phi \equiv \Delta[-L, L] \times [t_0, 1]$. Then the lowest and highest possible expected utilities over the set of priors are the ones based on degenerate distributions that assign the entire probability mass to the values of (M, t) resulting in the lowest and highest second-stage equilibrium profits. In this way, the two specifications coincide⁴.

Another objective specification that may be used to resolve the considered form of ambiguity is the minmax regret formulation. Define:

$$x_i^{br} = \arg \max_{x_i} \pi_i^*(x_i, x_{-i}, M, t)$$

and let the payoffs in the reduced first-stage game be given by:

$$\Pi_i^r(x_1, x_2) = - \max_{(M,t) \in S} \left[\pi_i^*(x_i^{br}, x_{-i}, M, t) - \pi_i^*(x_i, x_{-i}, M, t) \right]$$

In what follows, both alternative payoff specifications will be considered and compared, before being subject to endogenous objective choice in Section 4.

⁴Function Π_i^α may be also related to the dual non-expected utility formulation by Yaari [1987] under unique prior information. Suppose there are only two possible outcomes of the uncertainty: (M^1, t^1) and (M^2, t^2) , occurring with equal probabilities, and the probability weighting function $\varphi(p)$ assigns a value of $(1 - \alpha)$ to $p \leq 1/2$ and a value of α thereafter. Then the higher rank (larger profit) outcome (M^i, t^i) will receive a weight $\pi(1/2) = 1 - \alpha$, while the lowest-profit outcome will be assigned a weight $1 - \pi(1/2) = \alpha$, just as in the present specification. A related analogy between pessimism under ambiguity and risk-aversion under probabilistic information is covered in Section 6.

3 Results: Exogenous Objective

The second-stage unique Nash Equilibrium profits are exactly the same as the ones derived in MZ, i.e.:

$$\pi_i^*(x_1, x_2, M, t) = \begin{cases} t(x_2 - x_1) \left[1 + 2(-1)^i (M - \bar{x}) \right] & \text{for } (-1)^i (M - \bar{x}) \geq 3/2 \\ t(x_2 - x_1) \left[3(-1)^i + 2(M - \bar{x}) \right]^2 / 18 & \text{for } (M - \bar{x}) \in (-3/2, 3/2) \\ 0 & \text{otherwise} \end{cases}$$

where $\bar{x} = (x_1 + x_2)/2$. The first (topmost) segment of the above piecewise function corresponds to firm i capturing the entire market (later referred to as “monopolistic equilibrium”), while the middle segment is where firm i shares the market with the rival (“competitive equilibrium”). It is immediately clear that $\pi_i^*(x_1, x_2, M, t)$ is increasing in t and straightforward to verify that it is also decreasing in M for $i = 1$ and increasing in M for $i = 2$. In other words, the second stage equilibrium profit of the firm located on the left declines as the customers are located further and further to the right. Similarly for the firm located on the right when the consumer preferences shift leftward. Consequently, the α -maximin objective function simplifies to:

$$\Pi_i^\alpha(x_1, x_2) = \alpha \pi_i^*(x_1, x_2, -L[-1]^i, t_0) + (1 - \alpha) \pi_i^*(x_1, x_2, L[-1]^i, 1)$$

Note that each player can ensure a positive profit in any state of nature, by locating at $x_i = -x_{-i}$, so that $\bar{x} = 0$ and $(M - \bar{x}) \in (-3/2, 3/2)$ for all $M \in [-L, L]$ (recall $L \leq 1/2$ by assumption), i.e. the second stage equilibrium is always competitive.

The following intermediate result will be useful in the identification of equilibrium locations.

Lemma 3.1 *Suppose locations are such that $x_1 \leq 0 \leq x_2$ and $(M - \bar{x}) \notin (-3/2, 3/2)$ for some $M \in [-L, L]$. Then it is possible for the player who then monopolizes the market to change location in a way that improves her profits for all $(M, t) \in S$.*

The above result simplifies the derivation of equilibria in the reduced first-stage game, making it possible to focus on locations resulting in competitive equilibria at stage two. Note that the $x_1 \leq 0 \leq x_2$ assumption must also be satisfied by any location equilibrium, since otherwise the firm with a smaller hinterland could relocate to the other side of the competitor and increase its payoff. For instance, if $x_1 = 1/2$ and $x_2 = 3/2$, then firm 2 can move to $-1/2$, which means product differentiation is the same as before, but for every $M \in [-L, L]$ the second-stage equilibrium profits will exceed those that would have been achieved for the old location (3/2) and a symmetric demand location $M' = -M$ (see the profit specification at the beginning of this section and the proof of Proposition 3.1 for

more details). This reasoning applies to both the α -maximin and minmax regret objective specifications, so that, in the former case, payoffs in the location equilibrium must be given by:

$$\Pi_i^{cc}(x_1, x_2) = (x_2 - x_1) \left\{ \alpha t_0 \left[(-1)^i (3 - 2L) - 2\bar{x} \right]^2 + (1 - \alpha) \left[(-1)^i (3 + 2L) - 2\bar{x} \right]^2 \right\} / 18$$

Leading to the following result.

Proposition 3.1 *Consider a variant of the Hotelling duopoly game in which the joint distribution of transport costs and the median of the (uniform) consumer preferences is unknown with support on $[t_0, 1] \times [-L, L]$, where $t_0 \in (0, 1]$ and $L \in [0, \frac{1}{2}]$. Suppose the firms choose locations so as to maximize the α -maximin objective function $\Pi_i^\alpha(x_1, x_2)$, given a common degree of pessimism $\alpha \in [0, 1]$. Then the unique equilibrium locations are:*

$$x_1^* = -x_2^* = \frac{(3 + 2L)^2 (1 - \alpha) + (3 - 2L)^2 t_0 \alpha}{4 [2L([t_0 + 1] \alpha - 1) - 3 - 3(t_0 - 1)\alpha]} \quad (1)$$

We now turn to the comparative statics of the above solution.

Corollary 3.1 *The equilibrium product differentiation $\Delta^* = 2x_2^*$ is decreasing in the degree of pessimism (α), as well as in the degree of uncertainty about transport costs (t_0).*

In addition, the effect of a change in L (which measures the degree of demand location uncertainty) on Δ^* reveals a crucial difference from the MZ results.

Corollary 3.2 *When the players are sufficiently pessimistic, the equilibrium product differentiation is decreasing in demand location uncertainty. More specifically:*

$$\partial \Delta^* / \partial L < 0 \Leftrightarrow \alpha > \hat{\alpha} = (3 + 2L)^2 / [(3 + 2L)^2 + (3 - 2L)^2 t_0]$$

where $\hat{\alpha}$ is decreasing in t_0 and increasing in L , so that for $\alpha > 4 / (4 + t_0)$ the equilibrium product differentiation is decreasing in L over the whole range of the parameter.

The effect of the parameters on second stage equilibrium prices and profits is a consequence of their impact on Δ^* . As the price equilibria resulting from x_i^* are always competitive, we have:

$$\begin{aligned} \pi_i^*(x_1^*, x_2^*, M, t) &= \Delta^* t \left[3(-1)^i + 2M \right]^2 / 18 \\ p_i^*(x_1^*, x_2^*, M, t) &= \Delta^* t \left[3 + 2(-1)^i M \right] / 3 \end{aligned}$$

and since $M \in [-1/2, 1/2]$, the effect of α , t_0 and L on π_i^* and p_i^* has the same sign as their effect on Δ^* .

This shows that ambiguity attitudes determine the way in which the firms respond to changes in the spectrum of possible demand variations when their exact probability distribution is unknown. On the one hand, it follows from condition $\alpha > \hat{\alpha}$ that firms taking an optimistic approach ($\alpha < 1/[1 + t_0]$) always respond to an increase of uncertainty by venturing further away from one another, which is consistent with MZ. On the other hand, when the duopolists are moderately pessimistic relative to the minimum transport costs ($\alpha \in [1/(1 + t_0), 4/(4 + t_0)]$), they initially decrease product differentiation when L increases, but reverse this tendency when uncertainty becomes sufficiently large. Finally, highly pessimistic firms ($\alpha > 4/[4 + t_0]$) always locate closer together when uncertainty increases, in contrast with MZ.

What drives the results is the fact that the players' α -maxmin approach prescribes them to resolve the uncertainty they face in different ways, in the sense that they proceed 'as if' they each had a different prior over M and were aiming to maximize the expected value of second-stage profits. Specifically, the player located on the left effectively assumes that the probability of a 'worst-case scenario', in which the demand is as far to the right as possible and transport costs are at their lowest, is equal to α . When taking a pessimistic approach, the player will then want to locate relatively far to the right and close to the competitor, thereby improving her strategic position (and the resulting Nash Equilibrium profits) in the lowest profit case of $M = L$. This will occur at the cost of losing some of the strategic advantage in case of a favorable demand realization, but this has low-priority when α is high. Crucially, the other player associates the same probability α with an opposite market scenario ($M = -L$), so that, when pessimistic, she will want to locate relatively far to the left and will mirror player 1's shift towards the counterpart, rather than respond by moving away in order to relax the resulting second-stage price competition⁵. This is also interesting in the context of the 'certainty' Hotelling game, which was analyzed for various, not necessarily symmetric, customer distributions (see, for example Anderson et al. [1997], Meagher et al. [2008]), but always based on the firms having exactly the same expectations regarding the distribution of customers across the space of tastes and possible states of nature. The present paper demonstrates that, starting from a common degree of ignorance, the firms may end up acting 'as if' they maximized expected profits subject to non-identical priors.

⁵A notable caveat is that the players may not choose to re-locate to the other side of the competitor despite acting 'as if' the majority of consumers were bound to be located there. This is because doing so would cause the former lowest-profit demand realisation to become the highest-profit one and *vice versa*, i.e. the players would effectively switch their beliefs when switching sides.

In order to further understand the effect of L on the equilibrium locations, it is helpful to observe that for $(M - \bar{x}) \in (-3/2, 3/2)$ we have:

$$\partial\pi_1^*(x_1, x_2, M, t)/\partial x_1 = t(3 - 2M + x_1 + x_2)(2M - 3 - 3x_1 + x_2)/18$$

which in turn is concave in M . In other words, as the demand shifts more and more to the right, the marginal gains from moving in the same direction increase by less and less. Consequently, when L increases, any additional benefits from re-locating rightward which this brings about in the lowest profit scenario are smaller than the corresponding additional losses in the $M = -L$ case, the more so the larger the value of L . Hence, any gains from reducing product differentiation due to an increase of uncertainty will be outweighed by losses, unless L is sufficiently small relative to α , i.e. the losses are small relative to the importance of gains for the player's decision variable. For this reason, the players' tendency to differentiate their products is weakened when uncertainty increases only as long as it is not too big, which could mean that it never happens ($\alpha < 1/[1 + t_0]$) or that it is the case for the entire range of $L \in [0, 1/2]$ ($\alpha > 4/[4 + t_0]$).

It is interesting to observe that t_0 , reflecting the degree of uncertainty about transportation costs, affects the equilibrium locations, despite the fact that the corresponding transport cost parameter t has no effect on location decisions under certainty. This was possibly why the potential role of uncertainty about transport costs (or, in general, about the price elasticity of consumer demand) has been ignored by the relevant literature (MZ assume $t = 1$). On the one hand, this seems reasonable, because if t does not affect location decisions under certainty, then it should not matter that the firms do not know its exact value, since the optimal choice will be the same regardless of what it is. On the other hand, it overlooks the potential interaction between two types of uncertainty: about the customers' locations and about the transport costs that they incur. In particular, if a certain realization of consumer demand usually coincides with low transportation costs, then the resulting second-stage price competition is fierce and the equilibrium profits are low. Thus, a firm may choose not to locate in a way that would be advantageous in those circumstances if that means being further away from demand realizations associated with higher transportation costs and hence potentially more profitable. In a sense, locating under uncertainty is similar to designing a product to be sold in distinct markets, characterized by different consumer preferences and various degrees of price competition. It is therefore natural for the firms to target those of them where the consumers care more about the characteristics of the product than about its price, i.e. the ones which are less competitive.

In the current α -maxmin framework, the lowest-profit outcome entails transportation costs t_0 , so that a reduction of this parameter would make the worst-case scenario even more of a threat. Despite that, the firms are less determined to insure against it by staying close together, because they would need to sacrifice more in the optimistic scenario in order to improve their situation in the pessimistic one. Note that uncertainty about the intensity of price competition effectively makes the players more optimistic, by decreasing the weight

associated with the worst-case scenario. Recall $\pi_i^*(x_1, x_2, M, t) = t \times \pi_i^*(x_1, x_2, M, 1)$. As a result, it dampens the negative effect of pessimism on strategic product differentiation, leading to the observed tradeoff between α and t_0 . This is interesting, because the fact that re-scaling the transport costs fails to affect location choices in the classic Hotelling framework is somewhat paradoxical. In contrast, the current model shows that firms facing a possibility of lower transportation costs are more likely to venture out into more distant areas, relaxing price competition.

Overall, ambiguity attitudes cause more variety in the players' behaviour than the characteristics of the common prior in MZ. In the latter model, the equilibrium locations are also symmetric, with product differentiation $\Delta_{MZ}^* = 3/2 + 2\sigma^2/3$, σ^2 being the variance of the distribution of M . Because of the restriction on the support of this distribution ($L \in [0, 1/2]$), the maximum possible variance is $1/4$, and hence Δ_{MZ}^* ranges from $1\frac{1}{2}$ (certainty) to $1\frac{2}{3}$ (maximum uncertainty). In contrast, in the current model we have $\Delta^* = 3/2 + L$ for $\alpha = 0$ and $\Delta^* = 3/2 - L$ for $\alpha = 1$, i.e. product differentiation ranging from 1 to 2 depending on the size of the uncertainty.

Minmax Regret Objective

It is also interesting to see how the results would change when adopting the minmax regret objective instead of the α -maximin. To this end, observe first that the regret associated with a particular realization of the uncertainty $(M, t) \in S$ can be written as:

$$\pi_i^*(x_i^{br}, x_{-i}, M, t) - \pi_i^*(x_i, x_{-i}, M, t) = t \left[\pi_i^*(x_i^{br}, x_{-i}, M, 1) - \pi_i^*(x_i, x_{-i}, M, 1) \right]$$

and is hence always maximized at $t = 1$, making it possible at this point to assume $t_0 = 1$ without loss of generality.

Lemma 3.1 also makes it possible to once again restrict to competitive equilibria at stage 2. Otherwise, a player could improve her payoffs for every $M \in [-L, L]$, thereby also reducing the regret for every outcome of the uncertainty. This leads to the following result:

Proposition 3.2 *Consider a variant of the Hotelling duopoly game in which the probability distribution of the median of the (uniform) consumer preferences is unknown with support on $[-L, L]$, where $L \in [0, \frac{1}{2}]$. Suppose the firms choose locations so as to maximize the minmax regret objective function $\Pi_i^r(x_1, x_2)$. Then the unique equilibrium locations are:*

$$x_1^* = -x_2^* = \left(\sqrt{729 - 192L^2} - 45 \right) / 24 \quad (2)$$

As in the α -maximin case, it is useful to conduct the comparative statics of L .

Corollary 3.3 *Given locations (2), the equilibrium product differentiation is increasing in the degree of demand location uncertainty.*

In other words, using the minmax regret objective instead of the α -maximin reproduces the MZ comparative statics results. This is because, in order to minimize the highest possible regret, each player wants to locate in a way that makes the regret resulting from $M = -L$ equal to that associated with $M = L$ (see Proof). When L increases, the fact that $(-1)^i (\partial^3 \pi_i^* / \partial x_i \partial M^2) > 0$ implies that more opportunities are lost under the highest-profit demand (by not re-locating towards it), than in the lowest-profit case. Consequently, the regret associated with the worst-case scenario increases by less than that in the best-case one, making it necessary to move away from the rival in order to restore parity.

The demonstrated discrepancy between the α -maximin and minmax regret objective specifications (with respect to the role of demand uncertainty) makes it essential to investigate which objective function is more likely to become prevalent under the present conditions. This is studied in the following section.

4 Results: Endogenous Objective

We now turn to consider the possibility of the firms being able to commit to a chosen stage-two objective function, for instance, by appointing a cautious manager / CEO or by deliberately putting themselves in a position where losing customers due to a sudden change in preferences could mean bankruptcy, thereby making it essential to take the necessary precautions with the worst-case scenario in mind.

Suppose the choice of objective occurs at ‘stage zero’, i.e. before the actual location-then-price game begins. It will not matter for the outcome whether the firms choose their objective functions sequentially or simultaneously, so long as their decisions are known before setting locations. Firms are assumed to select their objective out of the α -maximin class (where $\alpha_i \in [0, 1]$ stands for the degree of pessimism of player i), in addition to the minmax regret specification.

Two restrictions will be imposed in order to simplify the analysis of this three-stage game. Firstly, since uncertainty about transport costs influences locations by effectively making an α -maximin player more optimistic, and has no effect under the minmax regret scheme, it will be assumed that $t_0 = 1$. Secondly, we will set $L = 1/2$, i.e. maximize the demand location uncertainty, in order to best distinguish between different objective functions (note that the equilibrium locations in Proposition 3.1 and Proposition 3.2 converge to $x_i^* = (-1)^i 3/4$ when uncertainty vanishes). This leads to the following result:

Proposition 4.1 *For any given objective choice of the counterpart, selecting an α -maximin objective function with $\alpha_i = 1$ results in the second-stage equilibrium profits being strictly higher for every $M \in [-1/2, 1/2]$ than those resulting from any other objective choice. Consequently, in equilibrium the duopolists choose $\alpha_1 = \alpha_2 = 1$.*

The result is due to the fact that when one player becomes more pessimistic, both shift towards the other player’s end of the market, with the player who changed her attitude

shifting more than the counterpart (see Proof, as well as Figure 1). This immediately




Figure 1: Iso-location curves of Player 1 and 2 (dashed), depending on the respective degrees of pessimism α_i . When one firm becomes more pessimistic, both shift towards the other firm's end of the market, but product differentiation decreases.

implies that the player whose attitude remains the same becomes worse off for all values of M . Most interestingly, the other player can improve her profits at the counterpart's expense, despite the reduced product differentiation (inducing more price competition) and regardless of the outcome of the uncertainty. Thus, $\alpha_i = 1$ is the only rational choice of objective at stage zero.

This suggests pessimism could serve as a way of strategic deterrence, discouraging the competitor from targeting one's market niche and instead making him withdraw into his own hinterland. However, a similar motive on behalf of the rival generates a Prisoner's Dilemma situation, as it was shown in the previous section that when both players become more pessimistic, they locate closer together and so earn less for all demand realizations. In this sense, the firms' self-imposed pessimism becomes a self-fulfilling prophecy.

Alternatively, rather than through conscious commitment at stage zero, the approach based on concentrating fully on the worst-case scenario ($\alpha_1 = \alpha_2 = 1$) could become prevalent *via* gradual elimination of underperforming, overly optimistic firms in repeated two-stage interactions. In any case, the contrasts between current results and the existing literature (see Corollary 3.2), should continue to hold under endogenous objective formulation. Indeed, given the predominance of an extreme form of pessimism, product

differentiation decreases in demand location uncertainty over the entire range of the associated parameter. Some of the observable industrial behaviour that is in line with this result, and could hence be explained in terms of the present model, is reported and discussed in the following section.

5 Example Applications

5.1 Sports Betting

To begin with, it is necessary to explain how the sports betting industry matches the premises of the current model. Consider the decision problem of a bookmaker, setting the odds on a sporting event with two possible outcomes, e.g. the team to win the football World Cup final. This is effectively a one-variable problem - whenever a bookmaker increases the odds on Team 1 to win, it decreases those on Team 2, so as to maintain a given ‘house edge’. More specifically, let o^j denote the winnings paid out per unit of currency bet on the j -th outcome in the event of it being realized, and let e represent the house edge. Next, consider a bettor who fully hedges her bets, i.e. bets a proportion β of her money on Team 1, and the remainder $(1 - \beta)$ on Team 2 in such a way that:

$$\overbrace{(1 + o^1) \beta}^{\text{Team 1 wins}} = \overbrace{(1 + o^2) (1 - \beta)}^{\text{Team 2 wins}}$$

The bookmaker can ensure that this would result in the bettor losing a fraction e of her money, by setting:

$$o^2 = [1 - e(2 + o^1)] / (e + o^1)$$

and making the odds on Team 2 a function of o^1 , which becomes the actual (sole) decision variable.

Suppose then the bettors’ estimates of the likelihood of Team 1 winning the game follow a certain distribution with support on the $[0, 1]$ interval. Everyone to the left of a certain point $\tilde{x}(o^1, o^2(o^1))$ will choose to bet on Team 2, while everyone else will bet on the rival team. The location of $\tilde{x}(o^1, o^2(o^1))$ will be further to the left when the relative odds on Team 1 increase, and overall will range from 0 to 1 depending on what the odds are⁶. In this way, by setting their odds, bookmakers effectively ‘locate’ in the space of the bettors probability estimates. Even in case of more than two possible outcomes, it could be argued that, since the bulk of betting volume is typically placed on the ‘favourite’, the problem can be simplified to that of setting the odds on the most likely (and popular) outcome, relative to those offered on the ‘underdogs’.

⁶This should happen regardless of whether bettors care only about their expected payoff, or exhibit risk-seeking tendencies.

The distribution of bettors' probability estimates often fluctuates according to patterns impossible for the bookmakers to elicit, mostly due to the continuing inflow of new information, interpreted by the bettors in ways difficult to predict. Thus, odds must be set under ambiguity, and only then it is possible to acquire more information by observing the actual distribution of bets placed on different outcomes.

The bookmakers' traditional objective in those difficult circumstances is to 'balance the books', by setting the odds in such a way that they will earn a profit regardless of what happens next, which makes it necessary to focus on the worst-case scenario, in which an outcome attracts an excessive volume of bets and is subsequently realized. This pessimistic approach may be supported by placing limits on the amounts individual players are allowed to bet on each outcome (thereby restricting the total amounts wagered), based on observing the odds set by other bookmakers in the first stage of the competition. These are important, because in the age of on-line betting a significant proportion of players are registered with several bookmakers and use odds comparison websites to select the one with the highest price on their preferred outcome. In this way, the bookmakers' decision problems are interlinked.

In general, product differentiation in the described market, measured by the dispersion of relative odds, tends to be quite high. As in the Hotelling model, the main reason for this is most likely that it allows the firms to avoid head-to-head competition, in this case in terms of the actual house edge. When relative odds tend to differ, different bookmakers will be best on different occasions in terms of the odds offered on *some* of the outcomes. In addition, especially for events with multiple outcomes, it is then very difficult for casual bettors to compare the house edge between various bookmakers.

In contrast, when product differentiation is small, i.e. relative odds are almost (or exactly) the same, then a bookmaker with a smaller house edge will be able to offer better odds on *all* outcomes and on *every* occasion. This creates pressure to reduce the edge, eventually driving it towards the 'fair' level. In this sense, a smaller degree of product differentiation is more competitive.

Nevertheless, an interesting phenomenon is that, particularly for racing events, the dispersion of odds greatly decreases shortly before the start of the competition⁷. This coincides with a significant increase of uncertainty about bettors' preferences in the corresponding time period, as evidenced by a largely increased price volatility in betting exchanges (see Smith et al. [2009]), possibly due to a lot of new information becoming available at this point.

The explanation might lie in the bookmakers' pessimistic attitude, and the fact that anyone choosing to offer different odds from the rest under increased uncertainty would face a particularly bad worst-case scenario. In particular, bettors might come to believe that the outcome on which the 'standout' bookmaker offers odds better than others is relatively

⁷This may be observed e.g. using a price comparison website, such as www.oddschecker.com, about twenty minutes before the start of any greyhound or horse race in the UK.

likely, making his bets unbalanced and profits low should the outcome in question be realized. In order to insure against this threat, it is better to offer the same as everybody else, because without bettors switching to and from other bookmakers in search of better odds any excessive volumes of bets resulting from a change in preferences will not be large enough to greatly unbalance the books. In other words, uncertainty then affects everyone in the same way and even an extreme pessimist sees no way of becoming disadvantaged.

An alternative explanation for the above phenomenon might be that perhaps the fraction of consumers using price comparison websites to select best odds greatly increases shortly before the start of a sporting event. This would mean that, in case of differentiated odds, each bookmaker ends up with an excessive volume of bets on the outcome on which it offers better odds than others, while little or no money is wagered on other outcomes. It would then be impossible to balance the books, unless odds are the same for every bookmaker (which is precisely what happens). However, while the data regarding the number of referrals from price comparison sites relative to the total volume of bets is not publicly available, there are no reasons to suggest such a tendency might occur. On the contrary, using comparison sites shortly before the start of the event does not seem very attractive for two reasons. Firstly, because of the delay between updating the odds by a bookmaker and a subsequent change in the comparison site, prices found in the latter constantly lag behind the actual ones at the time of increased preference (and price) volatility. Secondly, there may simply not be enough time to compare prices, wait to be re-directed to bookmaker, finally to log-in and place the wager, at the point when the event is about to start. Consequently, explaining the phenomenon in terms of the current model appears more convincing.

5.2 Mutual Funds

Another example is provided by the managed mutual funds' market. In this case, the Hotelling line may approximate a portfolio's position on the hypothetical Markowitz efficient frontier - from safe money market investments to risky equity funds. Once again, less product differentiation should prove more competitive, as it induces a Bertrand-like rivalry in terms of the funds' fees.

The more interesting it is to observe a tendency of fund managers to make their products more alike in the period following the financial crisis of 2008, than in a similar period preceding the market disturbances. The trend is evidenced by the increased average correlation between the daily returns of the fifteen most popular⁸ actively managed US funds (i.e. excluding 'index tracker' funds and passive money market funds). This was up from 0.75 in the period from June 2006 to January 2008, to 0.94 between June 2009 and January 2011⁹. See also Figure 2 for the comparison of the respective mean return and volatility statistics.

⁸According to www.marketwatch.com

⁹Data obtained via the Wolfram Alpha project, see online Appendix @13 for the details.




Figure 2: The means and standard deviations of the daily returns on the fifteen most popular actively managed US mutual funds, pre-2008 (circles) and post-2008 (squares).

The explanation might lie in the fact that the financial crisis forced financial service providers to revise their views about the range of individual investor behaviour they might reasonably expect and should be prepared to. In particular, the sharp decline of stock prices throughout 2008 and the beginning of 2009 resulted in a strong shift of consumer preferences towards relatively safe assets. This put pressure on mutual funds, which had to cope with an unprecedented scale of capital withdrawal requests, threatening the funds' liquidity.

The market resumed a pattern of relatively steady growth from mid-2009. Nevertheless, the recent events had to be accounted for in stress test simulations, evaluating the performance of portfolios in hypothetical extreme situations. These tests were either imposed on fund issuing institutions by government regulation, or were conducted in the funds' own capacity, as a way of reassuring investors about the safety of their assets. In addition, in the light of the present model, performing and satisfying a conservative stress test may be interpreted as a signal sent out to the competitors of one's commitment to a pessimistic, worst-case oriented approach.

In any case, such tests now had to acknowledge a possibility of market turbulences causing a sudden shift of consumer preferences away from the portfolio in question, making the worst-case scenario even worse than before and increasing uncertainty. In line with the current model, this resulted in a decrease of product differentiation, in the sense that the daily returns on different funds started moving more closely together in time, leading to the portfolios in question being clustered on the efficient frontier, with similar mean return and volatility values. Based on the present discussion, this could be seen as a way of

insurance against losing customers to other funds, which was less likely to happen if the alternatives were not much different from one's own product. However, it came at a cost of more intense competition in terms of the funds' fees, at a time when the ratios of fixed management costs to funds' assets have soared due to a decrease in the latter. Despite this increase in average costs, the fees of actively managed US funds have been on average slightly lower during the second of the two time periods considered above (June 2009 to January 2011)¹⁰.

6 Further Discussion

6.1 Social Welfare Analysis

It is interesting to consider the welfare implications of changes in the firms' objective, as well as the degree of uncertainty. Because of the assumption of inelastic demand, welfare is evaluated via (minus) the total transportation cost incurred by the consumers under certainty, or its expected value with respect to the probability distribution of the random shock in the MZ specification. Although that distribution is unknown to the firms in the present context, its exact form is largely inconsequential for comparative statics. To see why, observe the following result.

Proposition 6.1 *Given symmetric locations $x_2 = -x_1 = \Delta/2$, the ex-post total transportation cost resulting from every $(M, t) \in [-1/2, 1/2] \times (0, 1]$ is increasing in Δ , so long as $\Delta > 7/9$.*

Indeed, Proposition 3.1 and Proposition 3.2 imply that the equilibrium product differentiation (Δ^*) is never smaller than 1, whichever objective function is used, i.e. $\Delta > 7/9$ in equilibrium. In addition, given an α -maximin objective, Δ^* is decreasing in α (recall Corollary 3.1), and is smaller for $\alpha = 1$ than when firms use minmax regret. Thus, Proposition 6.1 implies that when players become more pessimistic, the ex-post social welfare improves for any realization of the uncertainty. Equivalently, whatever the actual distribution of (M, t) , the equilibrium expected total transportation cost, calculated with respect to that distribution, is smaller for $\alpha = 1$ than for any other objective. Clearly, this also implies that $\alpha = 1$ would still be best for social welfare if the latter was measured using the α -maximin or minmax regret criteria, instead of the expected value.

Suppose then that a social planner can influence the way the firms' objectives are set, for instance, by making mutual fund stress tests mandatory in the example in Section 5.2. The above argument then suggests that the planner would find the $\alpha = 1$ objective most desirable. This would be the case regardless of whether the alternative objective

¹⁰Source: *Trends in Mutual Fund Fees and Expenses in: Oversight of the Mutual Fund Industry: Ensuring Market Stability and Investor Confidence*, report of Investment Company Institute (www.ici.org) for the US House of Representatives Committee on Financial Services (June 2011)

specifications are evaluated based on a prior probability distribution (known to the planner, but not the firms), or one of the criteria available to the duopolists as a way of resolving ambiguity.

In fact, Proposition 4.1 established that $\alpha = 1$ is also the objective specification adopted by both firms in the unique equilibrium of the endogenous objective choice game considered in Section 4. Thus, the above discussion may be summarized in the following way.

Corollary 6.1 *Endogenous objective choice results in the implementation of the socially-optimal objective function.*

Proposition 6.1 is equally useful in investigating the welfare effects of changes in the amount of information available to the firms, as captured by the parameters L and t_0 . In other words, suppose that the actual distribution of (M, t) remains unchanged and unknown to the firms in its exact form, but what changes is the firms' knowledge of the support of the distribution (and the resulting degree of uncertainty), which they believe to be $[-L, L] \times (t_0, 1]$. Corollaries 3.1 and 3.2 report how this affects the equilibrium level of product differentiation, while Proposition 6.1 reveals how the latter affects welfare. Thus, we have:

Corollary 6.2 *When the players are sufficiently pessimistic, in the sense that*

$$\alpha > \hat{\alpha} = (3 + 2L)^2 / [(3 + 2L)^2 + (3 - 2L)^2 t_0]$$

social welfare improves as demand location uncertainty increases (L goes up). Otherwise, welfare declines with L . When uncertainty about transport costs increases (t_0 goes down), welfare always deteriorates.

Once again, the results are in contrast with MZ, who find that, under probabilistic information, demand location uncertainty is detrimental to social welfare. Here, the opposite occurs when the players are sufficiently pessimistic, which was demonstrated to be the case in equilibrium under endogenous objective choice. This could have important policy implications - for instance, for the regulation of market research practices. A common argument invoked to justify the gathering of information about consumer preferences (e.g. recently via social-networking sites) is that firms that know more about their customers' tastes at the stage of product design will subsequently offer products better suited to those tastes. However, the present model demonstrates that, paradoxically, the opposite might be the case in the absence of exact probabilistic information about potential demand variations.

6.2 Relationship to Risk Aversion under Probabilistic Information

The driving force behind the results in Section 3 is the fact that pessimistic players care more about those states of nature in which they earn less. Since these are the ones in which

the demand shifts towards the competitor, such players choose locations that are further in the same direction.

It appears that an analogous tendency should occur under probabilistic information if players were maximizing the expected value of a *concave* function of second-stage equilibrium profits. This would make them similarly focused on improving their profits in those states of nature in which they are relatively low.

As an example, let us consider a particularly interesting situation in which firms are interested in maximizing a concave transformation of their initial profits, arising under a progressive corporate taxation scheme. More specifically, suppose the scheme entails a tax rebate on profits up to R , and a constant tax rate of T thereafter (as in many actual tax systems, e.g. that of Singapore). Suppose further that the model is otherwise the same as in Section 3, except it is known to the firms that M is either equal to L or $-L$, with equal probabilities, and that t is always equal to 1.

It follows from MZ that $T = 0$ would result in the equilibrium product differentiation being equal to $\Delta^h = 3/2 + 2L^2/3$, while its lowest possible value, $\Delta^l = 3/2 - L$, is where each player is entirely focused on the realization of M that is as far as possible in the direction of the rival (this can be obtained by setting $\alpha = 1$ in Proposition 3.1).

Proposition 6.2 *In the model described above, any degree of product differentiation Δ such that $\Delta \in (\Delta^l, \Delta^h)$ can be implemented in equilibrium by means of a tax schedule which entails:*

$$R^* = \Delta(2L - 3)^2/18, \quad T^* = 2(6\Delta - 4L^2 - 9) / [(2\Delta - 2L - 3)(3 + 2L)]$$

The fact that progressive taxation in the form of tax rebates may be used to reduce the excess product differentiation and improve welfare is interesting in its own right. It also makes it even more striking that the same optimal outcome of $\Delta = \Delta^l$ which the tax may (asymptotically) implement under probabilistic information, can also be achieved under ambiguity via endogenous choice of objective and without the need for government intervention.

Nevertheless, while it is possible to construct examples, such as the one above, in which risk aversion replicates the effect of pessimism under ambiguity, it should be mentioned that this is not generically the case. Risk aversion under probabilistic information constitutes a much broader and complex problem, where instead of clear-cut, general results, much is likely to depend on the exact specification of the players' utilities, as well as the precise nature of the probabilistic information at their disposal.

In particular, the difference between risk aversion and pessimism becomes apparent when considering the interaction between the two types of uncertainty discussed in the present paper. For instance, suppose we introduce uncertainty in the form of four possible states of nature, characterized by different combinations of $M \in \{-L, L\}$ and $t \in \{t_0, 1\}$. Suppose further the corresponding probabilities are such that it is equally likely to have $M = -L$ and $M = L$, but that $t = t_0$ is much more likely in the former case. If, in addition,

t_0 is sufficiently small, than Firm 1 may actually earn on average less when $M = -L$ than when $M = L$. Thus, the more risk averse it becomes, the more it will care about profits associated with $M = -L$ relative to those attained at $M = L$, and the further away from the competitor it will want to locate, possibly in excess of what would be optimal before introducing uncertainty. In contrast, if the probabilities of the four states of nature were unknown, then players who are more pessimistic would continue to locate closer together, where a sufficient degree of pessimism would ensure that product differentiation is lower than under certainty.

Most importantly, a model with risk-averse firms under probabilistic information is not a substitute for the present framework, since they are each applicable to different real-world situations. For instance, the fact that ambiguity is prevalent in financial markets and the betting industry is by now widely recognized and accounted for in the literature (see e.g. Rigotti and Shannon [2005] for an overview, and recall the examples discussed in Section 5). In such cases, it is unrealistic to expect that market participants are able to agree on a unique prior, and hence a model allowing for ambiguity is more relevant. In addition, the simplicity of the α -maximin and minmax regret criteria permits an analysis more thorough and tractable than one relying on risk-aversion whilst going beyond the toy model presented in this section.

Nevertheless, while a complete analysis of the impact of risk aversion under probabilistic information is beyond the scope of the present paper, it constitutes an interesting possibility for future research.

6.3 Robustness

6.3.1 Sequential Location Choice

The key results of the paper are robust to a change in the timing of moves such that one of the players (say, Firm 1) chooses her location first. To see this, suppose $t_0 = 1$ and $\alpha = 0$, so that the first-stage payoffs are:

$$\Pi_i^0(x_1, x_2) = \pi_i^*(x_1, x_2, L[-1]^i, 1)$$

Consequently, the best response of Firm 2 is $x_2^{br}(L) = (3 + 2L + x_1)/3$ (see Proof of Proposition 3.2), and the equilibrium location of Player 1 is:

$$x_1^* = \arg \max_{x_1} \pi_1^*(x_1, x_2^{br}(L), -L, 1) = 0 \Rightarrow x_2^* = (3 + 2L)/3$$

which implies that product differentiation increases with uncertainty. In contrast, $\alpha = 1$ would mean that the first-stage payoffs are:

$$\Pi_i^1(x_1, x_2) = \pi_i^*(x_1, x_2, -L[-1]^i, 1)$$

leading to:

$$x_1^* = \arg \max_{x_1} \pi_1^* \left(x_1, x_2^{br}(-L), L, 1 \right) = 0 \Rightarrow x_2^* = (3 - 2L) / 3$$

Thus, despite the occurrence of first mover's advantage, results of Section 3 continue to hold, in that pessimistic players locate closer together than optimistic ones, and further reduce product differentiation when uncertainty increases.

6.3.2 Pricing under Uncertainty

We may follow MZ in considering the possibility of a change of timing, so that *both* stages of the game are played before uncertainty is resolved. For the firms uncertain only about the demand's location (but not the transport costs) the model is then relatively straightforward. Since the price is the same for all demand realizations, the payoff associated with a particular set of locations and prices is equal to a firm's own price multiplied by the α -maxmin value of the corresponding demand, i.e. of the consumer mass located on the relevant side of the 'indifferent consumer'. Hence, each player acts 'as if' being involved in a certainty Hotelling game in which the customer distribution is given by a weighted average of the demand realization located as far as possible towards the rival and the one at the opposite extreme, with α and $1 - \alpha$ being the respective weights. Consequently, a pessimistic player will consider an increase of uncertainty in similar terms as a shift of a small mass of consumers into his own hinterland matched by an opposite shift of a larger mass of customers into the hinterland of the competitor. Naturally, such a change would persuade a sufficiently pessimistic firm to move towards the rival, i.e. more uncertainty would decrease product differentiation and prices, as in the model discussed in the previous section.

The situation is somewhat complicated with the introduction of uncertainty about transport costs. It turns out that this may induce multiplicity of price equilibria at stage two, thereby precluding a unique solution at the location stage.

Proposition 6.3 *Consider a variant of the game described in Section 3 in which prices are chosen before uncertainty is resolved so as to maximize the α -maxmin value of the resulting profits. Whenever the first-stage locations satisfy:*

$$x_1 + x_2 \in [-x_T, x_T], \quad x_T = (1 - 2\alpha) [(4\alpha - 2)L - 1] (1 - t_0) / (1 + t_0)$$

there exists a continuum of second-stage Nash Equilibria, in which the firms set identical prices. The lowest and highest possible equilibrium prices for any pair of locations both decrease as L increases, provided that $\alpha \in (1/2, 1]$. However, the effect is dampened by a fall in t_0 .

Apart from the complication associated with the existence of multiple equilibria for 'not too asymmetric' locations, the results are in line with those in Section 3. For pricing

decisions made upon learning the consumer preferences, an increase of uncertainty led pessimistic firms to reduce product differentiation, thereby indirectly decreasing the resulting equilibrium prices. In the present case, uncertainty causes a similar change in prices by affecting the pricing decisions directly. Either way, the increase in the intensity of price competition is dampened by a possibility of a more price responsive consumer demand (lower costs).

6.3.3 Alternative Transport Cost Specification

Finally, the results should hold for any convex cost specification that permits the existence of a unique second-stage equilibrium in prices and provides the incentives for strategic product differentiation at stage one (e.g. a power function, as studied by Economides [1986]). Those incentives could then be offset by the players' pessimistic approach, just as it happens in the current model for quadratic costs.

Nevertheless, as is pervasive in the Hotelling literature, results can be invalidated by introducing a cost formula that does not have the above desirable properties, e.g. the basic linear specification¹¹. In the present context, this would mean that, in addition to ambiguity about the distribution of demand fluctuations, players would be faced with strategic uncertainty about the counterpart's price choices if the first-stage locations were sufficiently close together.

The problem could be circumvented by assuming that players consider uncertainty about acts of nature and those of the competitor in a similar manner. For instance, the worst-case scenario could entail an unfavourable shift in consumer preferences combined with competitive pricing on behalf of the counterpart. However, since the motive for strategic product differentiation is absent under linear costs, it may then be altogether unnecessary to study how ambiguity paired with pessimism can work to offset it.

7 Concluding Remarks

The paper examined a variation of the spatial duopoly framework in which the firms are unaware of the exact distribution of potential demand fluctuations, but resolve the resulting ambiguity using either the α -maximin or minmax regret criterion. It was demonstrated that players who are sufficiently pessimistic (characterized by a high α) locate closer together when uncertainty about the demand's location increases, which contrasts the existing results on spatial competition under probabilistic information and could explain some of the observed industrial behaviour. In addition, uncertainty about the responsiveness of consumer demand to price changes (reflected by the scale of transport costs) was shown to dampen the above tendency, by effectively making the players more optimistic.

¹¹But see Krol [2011] for a study of how a finite consumer reservation price may alleviate the problem under probabilistic demand location uncertainty.

Intuitively, a large range of possible demand variations is potentially more harmful to a player when product differentiation is bigger. Hence, a pessimistic entrepreneur may want to insure against this threat and respond to an increase of uncertainty by making her offer more similar to that of the other firm, leaving less room for her product becoming disadvantaged. This in turn would induce the players to price more aggressively once uncertainty is resolved, making the market more competitive as a result.

Indeed, commitment to a highly pessimistic objective may become prevalent as a method of strategic deterrence, signalling one's competitive, worst-case oriented approach to the rival prior to location choice, thereby prompting him to give ground and withdraw into his own hinterland. In particular, it was demonstrated that when players can commit to a chosen objective function, the unique equilibrium exhibits an extremely pessimistic maximin approach, which dominates any other α -maximin or the minmax regret specification.

The firms are essentially in a Prisoner's Dilemma situation - while being more pessimistic makes one better off at the rival's expense regardless of the resolution of the uncertainty, both firms lose out if they both become more focused on the worst-case scenario. This is in contrast with existing literature on behavioural industrial organization, which suggests that entrepreneurial optimism should outperform a pessimistic approach, leading to less competition. Nevertheless, the same maximin attitudes which are so damaging to the firms' profits once implemented via endogenous objective choice, were also demonstrated to be best from the social perspective. Thus, lack of precise probabilistic information about consumer demand can, rather strikingly, ensure that products are better suited to the consumers' needs, leading to a competitive and socially desirable outcome.

Appendix

Some of the longer derivations are relegated to an online Wolfram Mathematica appendix (<http://tinyurl.com/3g97oyn>), where numbers preceded by ‘@’ represent a position therein where a particular statement is verified. A free Mathematica Reader software is available at: <http://www.wolfram.com/products/player/>, and a non-interactive PDF version of the on-line appendix is available at: <http://tinyurl.com/3jmbt7o>.

Proof of Lemma 3.1

Consider the following cases:

[1.] Player 1 captures the entire market for a range of values of M (which must include $M = -L$), while player 2 does the same for some other M (including $M = L$). This is impossible, since:

$$(L + \bar{x} \geq 3/2 \wedge L - \bar{x} \geq 3/2) \Leftrightarrow L - 3/2 \geq |\bar{x}|$$

which is false by the assumption of $L < 1/2$.

[2.] The same player i monopolizes the market for all $M \in [-L, L]$. Then the other player can locate at $x_{-i} = -x_i$ and ensure a positive profit in any state of nature.

[3.] Player 1 monopolizes the market for some M (including $M = -L$) and a competitive equilibrium ensues for all the other M (including $M = L$). We then have:

$$\frac{\partial \pi_1^*}{\partial x_1}(x_1, x_2, -L, t) = t(3 + 2L + x_1 - 3x_2)(3 + 2L - x_1 - x_2)$$

The last part of the above product is positive by the $L - \bar{x} > -3/2$ assumption (competitive requirement for $M = L$). The other part is negative, since the monopolistic requirement for $M = -L$ is $L + \bar{x} > 3/2 \Leftrightarrow x_2 > 3 - 2L - x_1$, which implies that:

$$3 + 2L + x_1 - 3x_2 < 4x_1 + 8L - 6$$

which in turn is less than $8L - 6 < 0$ (by $x_1 < 0$ and $L < 1/2$), so that we have:

$$\frac{\partial \pi_1^*}{\partial x_1}(x_1, x_2, -L, t) < 0$$

Payoffs $\pi_i^*(\cdot)$ are the same as in the unconstrained ‘certainty’ Hotelling model (see Lambertini [1994]) and are hence single-peaked, where a change in M will cause a shift of the optimal location in the same direction. Hence, if x_1 is to the left of the peak for $M = -L$, then it must also be beneficial to shift rightwards when $M > -L$, until the equilibria are competitive for all $M \in [-L, L]$. By symmetry, a reverse statement holds when player 2 monopolizes the market for $M = L$ and a competitive equilibrium follows for $M = -L$.

Proof of Proposition 3.1

Based on Lemma 3.1, the equilibrium first order conditions are:

$$\frac{\partial \Pi_1^{cc}}{\partial x_1} = [1 + (t_0 - 1)\alpha] [4L^2 + (3 + 3x_1 - x_2)(3 + x_1 + x_2)] - 4L(3 + 2x_1) [(t_0 + 1)\alpha - 1] = 0$$

$$\frac{\partial \Pi_2^{cc}}{\partial x_2} = [1 + (t_0 - 1)\alpha] [4L^2 - (3 + x_1 - 3x_2)(x_1 + x_2 - 3)] + 4L(2x_2 - 3) [(t_0 + 1)\alpha - 1] = 0$$

This has three possible solutions, two of which fail to satisfy the competitive equilibrium condition $(M - \bar{x}) \in (-3/2, 3/2)$ for $M \in \{-L, L\}$ (online appendix @1). The remaining solution is (1), which gives both players an α -maxmin profit equal to:

$$\Pi^* = \frac{[(3 + 2L)^2 (1 - \alpha) + (3 - 2L)^2 t_0 \alpha]^2}{36 [3 + 3(t_0 - 1)\alpha - 2L([t_0 + 1]\alpha - 1)]}$$

Indeed, it turns out that no unilateral deviation from x_i^* can result in a payoff higher than Π^* . Firstly, player 1 cannot deviate to a $x_1 \neq x_1^*$ such that $x_1 < x_2^*$ and monopolize the market for some $M \in [-L, L]$, since this requires:

$$M - (x_1 + x_2^*)/2 \leq -3/2 \Rightarrow -L - (x_1 + x_2^*)/2 \leq -3/2 \Leftrightarrow x_1 \geq 3 - 2L - x_2^* \Rightarrow x_2^* > 3/2 - L$$

which is impossible under the current parameter assumptions (@2). In other words, an optimal location x_1 to the left of x_2^* must be such that the resulting price equilibria are competitive for $M \in \{-L, L\}$.

Out of all such locations, x_1^* is best, as it can be shown that $\partial \Pi_1^{cc}(x_1, x_2^*)/\partial x_1$ is negative for $x_1 < x_1^*$ and positive for $x_1 > x_1^*$, as long as $M - (x_1 + x_2^*)/2 \in (-3/2, 3/2)$ for $M \in \{-L, L\}$, see (@3).

This means it is also impossible to benefit from deviating to a $x_1 > x_2^*$. To see this, observe that the game is symmetric, in the sense that $\Pi_2(x_1, x_2) = \Pi_1(-x_2, -x_1)$. As a consequence, for each $x_1 > x_2^*$ there exists a location $x_1' = x_2^* - (x_1 - x_2^*) < x_2^*$ giving a higher payoff, since:

$$\Pi_1(x_1', x_2^*) - \Pi_2(x_2^*, x_1) = \Pi_1(2x_2^* - x_1, x_2^*) - \Pi_1(-x_1, -x_2^*)$$

which means switching from x_1 to x_1' is equivalent to shifting the locations of both players rightward by the same distance of $2x_2^* > 0$. This in turn has the same effect as shifting the customer distribution leftward for locations fixed, thereby increasing the second-stage profit of player 1 for $M \in \{-L, L\}$. All in all, it is impossible for player 1 to gain by deviating from x_1^* , while the converse is true for player 2 (again, by symmetry of the game). ■

Proof of Proposition 3.2

Given a particular value of M , the locations that maximize the second stage competitive equilibrium profit subject to $M - \bar{x} \in (-3/2, 3/2)$ are (@4):

$$x_i^{br}(M) = \left[(-1)^i 3 + 2M + x_{-i} \right] / 3$$

The regret $r_i(x_i, x_{-i}, M)$ of player i associated with a given M is then the difference between the competitive profit that would have been obtained by locating at $x_i^{br}(M)$ and the one actually obtained, which evaluates to:

$$r_i(x_i, x_{-i}, M) = \left[12 + (-1)^i (8M - 3x_i - 5x_{-i}) \right] \left[(-1)^i 3 + 2M - 3x_i + x_{-i} \right]^2 / 486$$

Since the competitive payoffs $\pi_i^*(\cdot)$ are single-peaked (in own location) and $x_i^{br}(M)$ is increasing in M , we may further restrict the equilibrium identification to locations $x_i \in [x_i^{br}(-L), x_i^{br}(L)]$. Otherwise, one could improve the profit (and decrease regret) in any state of nature. In addition, it can be shown (@5) that $r_i(x_i, x_{-i}, M)$ decreases in M so long as $x_i^{br}(M) < x_i$. Once M is high enough for $x_i^{br}(M)$ to exceed x_i , any further increase of M increases the player's regret. Thus, we have:

$$\max_{M \in [-L, L]} r_i(x_i, x_{-i}, M) = \max \{ r_i(x_i, x_{-i}, -L), r_i(x_i, x_{-i}, L) \} \quad (3)$$

As x_i increases (while within $[x_i^{br}(-L), x_i^{br}(L)]$), $r_i(x_i, x_{-i}, -L)$ increases, but $r_i(x_i, x_{-i}, L)$ falls. In order to minimize (3), players then want to locate at x_i such that:

$$r_i(x_i, x_{-i}, -L) = r_i(x_i, x_{-i}, L) \quad (4)$$

And it can be shown (@6) that locations (2) are the only ones that satisfy both the above equality and the condition $x_i \in [x_i^{br}(-L), x_i^{br}(L)]$ for $i = 1, 2$. ■

Proof of Proposition 4.1

Let $x_i^*(o_i, o_{-i})$ denote the unique stage-one equilibrium location of player i , given that she and the counterpart have chosen respective objective functions $o_i, o_{-i} \in \{r, [0, 1]\}$ at stage zero, where r stands for minmax regret and numbers from the unit interval represent the corresponding degrees of pessimism. It should also be noted that uniqueness of equilibria holds 'up to symmetry'. For instance, there is an equilibrium in which $x_1^*(0, 1) = -\frac{5}{4}$ and $x_2^*(1, 0) = \frac{1}{4}$, and one in which $x_1^*(1, 0) = -\frac{1}{4}$ and $x_2^*(0, 1) = \frac{5}{4}$, i.e. the pessimistic player is always located closer towards the centre of the market and both receive the same payoffs regardless of which configuration is selected. This of course raises a certain coordination problem, which is, however, no different from the classic Hotelling case, since each player gets to choose between two alternative equilibrium locations. See Bester et al. [1996] for the related discussion.

Based on Lemma 3.1, one can complement the locations in Proposition 3.1 and Proposition 3.2 with those that would ensue for asymmetric objectives. These are, once again, obtained by assuming competitive profits and solving the associated FOC's, either $\partial\Pi_i^{\alpha_i}/\partial x_i = 0$ or $r_i(x_i, x_{-i}, -L) = r_i(x_i, x_{-i}, L)$, for x_1 and x_2 . The only solution satisfying the competitive equilibrium requirements is then tested for optimality (see @7). The proof will be split into four parts:

Part I

The $\alpha_i = 1$ objective outperforms any $\alpha_i < 1$, so long as the counterpart uses some $\alpha_{-i} \in [0, 1]$. To see this, observe first that we have:

$$\begin{aligned} \pi_1^*(x_1^*(\alpha_1, \alpha_2), x_2^*(\alpha_2, \alpha_1), M) &= \\ &= [x_2^*(\alpha_2, \alpha_1) - x_1^*(\alpha_1, \alpha_2)] [-3 + 2M - x_1^*(\alpha_1, \alpha_2) - x_2^*(\alpha_2, \alpha_1)]^2 / 18 \end{aligned}$$

differentiating with respect to α_1 gives a product $A \times B$, where:

$$A = [-3 + 2M - x_1^*(\alpha_1, \alpha_2) - x_2^*(\alpha_2, \alpha_1)] / 18 < 0 \Leftrightarrow M - \bar{x}^* \in (-3/2, 3/2)$$

and:

$$\begin{aligned} B &= \frac{\partial x_2^*(\alpha_2, \alpha_1)}{\partial \alpha_1} [-3 + 2M + x_1^*(\alpha_1, \alpha_2) - 3x_2^*(\alpha_2, \alpha_1)] - \\ &\quad - \frac{\partial x_1^*(\alpha_1, \alpha_2)}{\partial \alpha_1} [-3 + 2M + x_2^*(\alpha_2, \alpha_1) - 3x_1^*(\alpha_1, \alpha_2)] \end{aligned}$$

which is negative for all $M \in [-1/2, 1/2]$ if and only if it is negative for $M = -1/2$, because it can be shown (@8) that $\partial x_1^*(\alpha_1, \alpha_2) / \partial \alpha_1 > \partial x_2^*(\alpha_2, \alpha_1) / \partial \alpha_1 > 0$, i.e. when a player becomes more pessimistic, both shift towards the other player's end of the market, with the player who changed her attitude shifting more than the counterpart. This immediately implies that the player whose attitude remains the same becomes worse off for all values of M . For the other player to become better off, we need:

$$\frac{\partial x_2^*(\alpha_2, \alpha_1)}{\partial \alpha_1} [x_1^*(\alpha_1, \alpha_2) - 3x_2^*(\alpha_2, \alpha_1) - 4] < \frac{\partial x_1^*(\alpha_1, \alpha_2)}{\partial \alpha_1} [x_2^*(\alpha_2, \alpha_1) - 3x_1^*(\alpha_1, \alpha_2) - 4]$$

which can be shown (@9) to be the case for all $\alpha_1, \alpha_2 \in [0, 1]$. Thus, player 1 improves profits for every $M \in [-1/2, 1/2]$ whenever increasing her degree of pessimism. A converse of this argument holds for a change in the attitude of player 2 (by symmetry of the game).

Part II

The $\alpha_i = 1$ objective outperforms any $\alpha_i < 1$, so long as the counterpart uses $\alpha_{-i} = r$. Similarly to the previous case, it can be demonstrated (@10) that:

$$\partial x_1^*(\alpha_1, r) / \partial \alpha_1 > \partial x_2^*(r, \alpha_1) / \partial \alpha_1 > 0$$

in addition to:

$$\frac{\partial x_2^*(r, \alpha_1)}{\partial \alpha_1} [x_1^*(\alpha_1, r) - 3x_2^*(r, \alpha_1) - 4] < \frac{\partial x_1^*(\alpha_1, r)}{\partial \alpha_1} [x_2^*(r, \alpha_1) - 3x_1^*(\alpha_1, r) - 4]$$

Where a converse again holds by symmetry of the game.

Part III

The $\alpha_i = 1$ objective outperforms $\alpha_i = r$, so long as the counterpart uses some $\alpha_{-i} \in [0, 1]$. The difference between the respective equilibrium payoffs in state M is:

$$\pi_1^*(x_1^*(1, \alpha_2), x_2^*(\alpha_2, 1), M) - \pi_1^*(x_1^*(r, \alpha_2), x_2^*(\alpha_2, r), M)$$

and can be shown (@11) to attain a minimum at $M = 1/2, \alpha_2 = 1$, in which case it is still positive.

Part IV

The $\alpha_i = 1$ objective outperforms $\alpha_i = r$, so long as the counterpart uses $\alpha_{-i} = r$. In this case, the difference between the respective equilibrium payoffs in state M is:

$$\pi_1^*(x_1^*(1, r), x_2^*(r, 1), M) - \pi_1^*(x_1^*(r, r), x_2^*(r, r), M)$$

which is quadratic concave in M and attains positive values at $M \in \{-1/2, 1/2\}$ (@12), which implies that using $\alpha_i = 1$ is always better. ■

Proof of Proposition 6.1

Let $x_1 = -x_2 = -\Delta/2$. Since the price equilibrium is then competitive for all (M, t) , we have:

$$p_i^* \left([-1]^i \Delta/2, -[-1]^i \Delta/2, M, t \right) = \Delta \left(3 + [-1]^i 2M \right) t/3$$

so that Firm 1 serves all customers with locations in $[M - 1/2, 2M/3]$ and everyone else is served by Firm 2. Consequently, the total transport cost incurred by the customers in state (M, t) is:

$$\begin{aligned} TC &= \int_{M-1/2}^{2M/3} t(x + \Delta/2)^2 dx + \int_{2M/3}^{M+1/2} t(x - \Delta/2)^2 dx = \\ &= t [9\Delta^2 - \Delta(9 + 20M^2) + 36M^2 + 3] / 36 \end{aligned}$$

and we have:

$$\partial TC / \partial \Delta = t(18\Delta - 20M^2 - 9) / 36 > 0 \Leftrightarrow \Delta > 1/2 + (10/9)M^2 \leq 7/9 \Leftrightarrow M \leq 1/2$$

i.e. $\Delta > 7/9$ is sufficient to ensure that $\partial TC / \partial \Delta > 0$ for all $(M, t) \in [-1/2, 1/2] \times (0, 1]$. ■

Proof of Proposition 6.2

To begin with, note that so long as $T < 1$, the tax has no effect on the second-stage pricing decisions, since any price that maximizes the gross profit also maximizes the amount net of tax. Consequently, the second-stage equilibrium prices (and gross profits) remain the same and Lemma 3.1 again makes it possible to restrict to locations resulting in competitive equilibria.

Let $x_1^* = -x_2^* = -\Delta/2$ be the equilibrium locations, and consider a possible deviation by player 1 to $x_1 \neq x_1^*$ such that:

$$\pi_1^*(x_1, x_2^*, -L, 1) > R^* = \pi_1^*(x_1^*, x_2^*, L, 1) > \pi_1^*(x_1, x_2^*, L, 1)$$

this holds for all $x_1 < x_1^*$ and may hold for some $x_1 \in (x_1^*, x_2^*)$. The expected net profit is then:

$$\Pi_1(x_1, x_2^*) = [R^* + (1 - T^*)[\pi_1^*(x_1, x_2^*, -L, 1) - R^*] + \pi_1^*(x_1, x_2^*, L, 1)]/2$$

Differentiating with respect to x_1 yields:

$$\frac{\partial \Pi_1(x_1, x_2^*)}{\partial x_1} = \frac{L(\Delta + 2x_1)[3\Delta - \Delta^2 + 16L^2 + 6(\Delta - 3)x_1 - 36]}{18(3 + 2L)(3 - 2\Delta + 2L)}$$

which can be shown (@17) to be positive for $x_1 < x_1^*$ and negative for $x_1 > x_1^*$. Another possibility is that player 1 deviates to a $x_1 \in (x_1^*, x_2^*)$ such that:

$$\pi_1^*(x_1, x_2^*, -L, 1) > \pi_1^*(x_1, x_2^*, L, 1) > R^*$$

so that the expected net profit equals:

$$\Pi_1'(x_1, x_2^*) = R^* + (1 - T^*)([\pi_1^*(x_1, x_2^*, -L, 1) + \pi_1^*(x_1, x_2^*, L, 1)]/2 - R^*)$$

and we have:

$$\frac{\partial \Pi_1'(x_1, x_2^*)}{\partial x_1} = \frac{(2L - 3)(2\Delta + 2\Delta - 3)[\Delta^2 - 4\Delta x_1 - 16L^2 - 12(1 + x_1)(3 + x_1)]}{72(-3 + 2\Delta - 2L)(3 + 2L)}$$

which can be shown (@18) to be negative for all $x_1 > x_1^*$. By symmetry, the same is true for Player 2 and $x_2 \neq x_2^*$. ■

Proof of Proposition 6.3

Since lower costs (i.e. higher price elasticity of the demand for a firm's product) are better for the lower priced firm, while higher costs are better for the firm with a higher price, the α -maximin profits of player i are given by:

$$\alpha \left[-L + 1/2 - (-1)^i \tilde{x} \right] p_i + (1 - \alpha) \left[L + 1/2 - (-1)^i \hat{x} \right] p_i$$

where \tilde{x} is the location of the indifferent consumer in the worst case scenario, i.e. the value of x that solves:

$$p_1 + t(x_1 - x)^2 = p_2 + t(x_2 - x)^2$$

with $t = 1$ if $p_i < p_{-i}$ and $t = t_0$ otherwise. Similarly, \hat{x} is the best possible indifferent consumer location, obtained for $t = t_0$ if $p_i < p_{-i}$ and $t = 1$ otherwise. The profit function is concave in p_i for $\alpha > 1/2$, but its first derivative is discontinuous at $p_i = p_{-i}$, so that the best-response functions are:

$$BR_i = \begin{cases} (p_{-i} + p_i^0)/2 & \text{for } p_{-i} > p_i^0 \\ (p_{-i} + p_i^1)/2 & \text{for } p_{-i} < p_i^1 \\ p_{-i} & \text{otherwise} \end{cases}$$

where:

$$p_i^j = \frac{t_0(x_2 - x_1) \left[1 + (2 - 4\alpha)L - (-1)^{i+j}(x_1 + x_2) \right]}{(t_0)^j + (-1)^{j+1} \alpha (1 - t_0)}$$

In other words, for a range of the counterpart's prices $[p_i^1, p_i^0]$ each firm would choose to respond with an identical price. This may coincide with a similar range of prices on the competitor's behalf, resulting in a continuum of Nash Equilibria of the price-subgame associated with locations x_1, x_2 :

$$\{(p_1, p_2) : (\exists p^* \in [p_1^1, p_1^0] \cap [p_2^1, p_2^0]) (p_1 = p_2 = p^*)\}$$

It is easy to check that this multiplicity of equilibria occurs when the firm locations satisfy:

$$x_1 + x_2 \in [-x_T, x_T], \quad x_T = (1 - 2\alpha) [(4\alpha - 2)L - 1] (1 - t_0) / (1 + t_0) \quad (5)$$

More specifically, the range of equilibrium prices is:

$$\begin{cases} [p_1^1, p_2^0] & \text{for } x_1 + x_2 \in [-x_T, 0] \\ [p_2^1, p_1^0] & \text{for } x_1 + x_2 \in [0, x_T] \end{cases}$$

while for $x_1 + x_2 \notin [-x_T, x_T]$ there is a unique asymmetric equilibrium. Assuming that locations satisfy 5, it may then be observed that:

$$\forall i, j : \frac{\partial p_i^j}{\partial L} < 0 \Leftrightarrow \alpha \in (1/2, 1] \Leftrightarrow \forall i, j : \frac{\partial p_i^j}{\partial L \partial t_0} < 0$$

which completes the Proof. ■

References

- Aghion, P., M. P. Espinosa, and B. Jullien (1993): “Dynamic duopoly with learning through market experimentation.” *Economic Theory*, 3, 517–537.
- Anderson, S. P., J. K. Goeree, and R. Ramer (1997): “Location, location, location,” *Journal of Economic Theory*, 77, 102–127.
- Armstrong, M. and S. Huck (2010): “Behavioral economics as applied to firms: A primer,” *Competition Policy International Journal*, 6.
- Arrow, K. and L. Hurwicz (1972): “An optimality criterion for decision-making under ignorance,” in C. Carter and J. Ford, eds., *Uncertainty and Expectations in Economics*, Blackwell, 1–11.
- Balvers, R. and L. Szerb (1996): “Location in the hotelling duopoly model with demand uncertainty,” *European Economic Review*, 40, 1453–1461.
- Bergemann, D. and K. H. Schlag (2008a): “Pricing without priors,” *Journal of the European Economic Association*, 6, 560–569.
- Bergemann, D. and K. H. Schlag (2008b): “Robust monopoly pricing,” *SSRN eLibrary*.
- Bester, H., A. de Palma, W. Leininger, J. Thomas, and E.-L. von Thadden (1996): “A noncooperative analysis of hotelling’s location game,” *Games and Economic Behavior*, 12, 165 – 186.
- Casado-Izaga, J. F. (2000): “Location decisions: The role of uncertainty about consumer tastes,” *Journal of Economics*, 71, 31–46.
- Economides, N. (1986): “Minimal and maximal product differentiation in hotelling’s duopoly,” *Economics Letters*, 21, 67–71.
- Englmaier, F. (2010): “Managerial optimism and investment choice,” *Managerial and Decision Economics*, 31, 303–310.
- Fabiani, S., C. Loupias, F. Martins, and R. Sabbatini, eds. (2007): *Pricing Decisions in the Euro Area: How Firms Set Prices and Why*, Oxford University Press.
- Ghirardato, P., F. Maccheroni, and M. Marinacci (2004): “Differentiating ambiguity and ambiguity attitude,” *Journal of Economic Theory*, 118, 133–173.
- Harter, J. F. R. (1997): “Hotelling’s competition with demand location uncertainty,” *International Journal of Industrial Organization*, 15, 327–334.

- Heifetz, A., C. Shannon, and Y. Spiegel (2007a): “The dynamic evolution of preferences,” *Economic Theory*, 32, 251–286.
- Heifetz, A., C. Shannon, and Y. Spiegel (2007b): “What to maximize if you must,” *Journal of Economic Theory*, 133, 31 – 57.
- Krol, M. (2011): “On the existence and social optimality of equilibria in a hotelling game with uncertain demand and linear-quadratic costs,” *The B.E. Journal of Theoretical Economics*, Vol. 11 : Iss. 1 (Topics), Article 7.
- Kyle, A. S. and F. A. Wang (1997): “Speculation duopoly with agreement to disagree: Can overconfidence survive the market test?” *The Journal of Finance*, 52, pp. 2073–2090.
- Lambertini, L. (1994): “Equilibrium locations in the unconstrained hotelling game,” *Economic notes*, 23, 438–46.
- Meagher, K. J., E. G. Teo, and W. Wang (2008): “A duopoly location toolkit: Consumer densities which yield unique spatial duopoly equilibria,” *The B.E. Journal of Theoretical Economics*, 8.
- Meagher, K. J. and K. G. Zauner (2004): “Product differentiation and location decisions under demand uncertainty,” *Journal of Economic Theory*, 117, 201–216.
- Meagher, K. J. and K. G. Zauner (2005): “Location-then-price competition with uncertain consumer tastes,” *Economic Theory*, 25, 799–818.
- Osborne, M. J. and C. Pitchik (1987): “Equilibrium in hotelling’s model of spatial competition,” *Econometrica*, 55, 911–922.
- Rigotti, L. and C. Shannon (2005): “Uncertainty and risk in financial markets,” *Econometrica*, 73, 203–243.
- Savage, L. J. (1951): “The theory of statistical decision,” *Journal of the American Statistical Association*, 46, pp. 55–67.
- Smith, M. A., D. Paton, and L. V. Williams (2009): “Do bookmakers possess superior skills to bettors in predicting outcomes?” *Journal of Economic Behavior and Organization*, 71, 539 – 549.
- Yaari, M. E. (1987): “The dual theory of choice under risk,” *Econometrica*, 55, pp. 95–115.