

GAME -THEORETIC MODELLING
OF OLIGOPOLISTIC COMPETITION
UNDER UNCERTAINTY

A thesis submitted to The University of Manchester
for the degree of Doctor of Philosophy in the Faculty of Humanities

2011

Michał Król

School of Social Sciences (Economics)

Contents

1	Introduction and Literature Review	10
1.1	Introduction	10
1.2	Literature Review: Heterogeneous Goods	11
1.3	Literature Review: Homogeneous Goods	16
1.4	Outline of Discussion	22
I	Heterogeneous Goods	25
2	On the existence and social optimality of equilibria in a Hotelling game with uncertain demand and linear-quadratic costs	26
2.1	Introduction	26
2.2	The Model	28
2.3	Results	31
2.4	Concluding Remarks	37
3	Product differentiation decisions under ambiguous consumer demand and pessimistic expectations	46
3.1	Introduction	46
3.2	The Model	50
3.3	Results	51
3.4	Extensions	59
3.5	Concluding Remarks	65

II	Homogeneous Goods	67
4	Games as Social Conventions	68
4.1	Introduction	68
4.2	The concept	71
4.2.1	Definition and basic properties	71
4.2.2	Relationship to vN-M stable sets	72
4.3	Application: Cournot vs. Bertrand	74
4.3.1	The model	74
4.3.2	Results	76
4.4	Concluding Remarks	81
5	Conclusion	88
	Bibliography	92

List of Figures

1.1	Hotelling Model with mill pricing	13
1.2	Quantity-Price Model by Moreno and Ubeda	21
3.1	Iso-location curves of the duopolists given asymmetric degrees of pessimism	61

Abstract

Game-Theoretic Modelling of Oligopolistic Competition under Uncertainty

Michał Król - The University of Manchester (Doctor of Philosophy)

6th October 2011

This thesis reconsiders some of the most widely debated controversies in game theoretic modelling of oligopolistic competition, and proposes modifications of existing models and concepts demonstrating how these controversies can be addressed in the presence of uncertainty.

The first part of the discussion is concerned with competition by product design in heterogeneous goods' markets, as captured by the Hotelling framework. The most prominent difficulty here is the model's lack of robustness to changes in the transportation cost specification, and the fact that the only universally tractable quadratic formulation induces an implausible and socially undesirable 'maximum differentiation' outcome. The problem is addressed in Chapter Two, by considering a model with uncertain consumer demand and general linear-quadratic costs. It turns out that, for uncertainty big enough, the presence of a linear component in the cost function no longer rules out an analytical solution to the game. In particular, I characterize a subgame-perfect equilibrium in which the firms' locations converge to the socially efficient ones in the limit as uncertainty increases, regardless of the curvature of the cost function. Thus, the presence of substantial demand uncertainty makes the market more competitive, by reducing the excessive equilibrium product differentiation and the resulting prices.

In fact, Chapter Three demonstrates that this also holds when the firms do not know the exact distribution of consumer demand fluctuations, but resolve the resulting ambiguity using the Arrow and Hurwicz α -maxmin criterion. This is because firms that are sufficiently pessimistic, in the sense of assigning a large weight to the lowest profit scenario, locate closer together in equilibrium under uncertainty, where it is argued that operating on such a pessimistic premise could become prevalent via strategic commitment, elimination of underperforming firms or as a result of taxation.

This discussion is complemented by the second part of the thesis, which is concerned with competition in *homogeneous* goods' markets. In particular, Chapter Four reconsiders the so far unresolved discrepancy between the Cournot model of quantity competition and the alternative Bertrand price setting specification. To this end, I propose a model evaluation criterion, based on a recent generalization of the von Neumann-Morgenstern stable set concept. In particular, a restriction of the players' strategy sets is said to constitute a stable convention when no one has an incentive to unilaterally violate it, while faced with strategic uncertainty about the counterparts' exact choices within their restricted strategy sets. Applying the criterion to a simultaneous move quantity-price game reveals that Cournot competition is a stable convention when production costs are high relative to the number of firms and difficult to recover for unsold output. In contrast, Bertrand competition is never stable.

DECLARATION

No portion of the work referred to in the thesis has been submitted in support of an application for another degree or qualification of this or any other university or other institute of learning.

COPYRIGHT STATEMENT

- i. The author of this thesis (including any appendices and/or schedules to this thesis) owns certain copyright or related rights in it (the “Copyright”) and s/he has given The University of Manchester certain rights to use such Copyright, including for administrative purposes.
- ii. Copies of this thesis, either in full or in extracts and whether in hard or electronic copy, may be made only in accordance with the Copyright, Designs and Patents Act 1988 (as amended) and regulations issued under it or, where appropriate, in accordance with licensing agreements which the University has from time to time. This page must form part of any such copies made.
- iii. The ownership of certain Copyright, patents, designs, trade marks and other intellectual property (the “Intellectual Property”) and any reproductions of copyright works in the thesis, for example graphs and tables (“Reproductions”), which may be described in this thesis, may not be owned by the author and may be owned by third parties. Such Intellectual Property and Reproductions cannot and must not be made available for use without the prior written permission of the owner(s) of the relevant Intellectual Property and/or Reproductions.

Further information on the conditions under which disclosure, publication and commercialisation of this thesis, the Copyright and any Intellectual Property and/or Reproductions described in it may take place is available in the University IP Policy (see www.campus.manchester.ac.uk/medialibrary/policies/intellectual-property.pdf), in any relevant Thesis restriction declarations deposited in the University Library, The University Library’s regulations (see <http://www.manchester.ac.uk/library/aboutus/regulations>) and in The University’s policy on presentation of Theses.

*Magdzie i Feliszce, dzięki którym nawet nie zauważyłem,
kiedy ten doktorat się napisał.*

ACKNOWLEDGEMENTS

I wish to thank Professor Paul Madden, without whose endless and patient guidance and feedback this thesis would not have been possible. I am also grateful to Professor Igor Evstigneev, to whom I could always turn to when struggling with mathematics, and who was particularly helpful with Chapter Four. Finally, thanks are due to Professor Nicholas Yannelis, who was always happy to advise me on any research related issues.

All the remaining errors are due to my son Felix crawling onto the keyboard when I wasn't looking.

Introduction and Literature Review

1.1 Introduction

Using game theory to study economic phenomena, despite offering substantial explanatory power, often requires that significant simplifying restrictions are imposed on the agents' individual decision problems, to ensure that the complexity of ensuing strategic interactions does not render the model intractable. Even so, alternative ways of simplifying the problem at hand may at times be given equally plausible justifications. More often than not, a theorist then opts for the model specification which results in a determinate outcome, in the sense that applying the Nash solution concept, or one of its later refinements, identifies a unique equilibrium.

Unfortunately, this raises questions as to the robustness of any comparative statics results obtained in this way. If a model relies on an ad hoc configuration of assumptions without empirical foundations, any conclusions it leads to, however interesting, may be dismissed as a modelling artefact having little to do with economic reality.

In this thesis, I will reconsider some of the most intensely debated modelling dilemmas and controversies within the theory of oligopolistic competition¹. In particular, the first part of the work, comprising two chapters, focuses on competition in heterogeneous goods' markets, where firms use product differentiation as a strategic instrument. The discussion

¹Although the focus here is on strategic interactions between sellers, the discussion could be, in large part, reinterpreted as a study of (labour market) oligopsony. See e.g. Bhaskar and To [2003] or Kaas and Madden [2010].

is set within the classic Hotelling framework and revolves around the sensitivity of results to the assumptions of quadratic vs. linear transportation costs and mill pricing vs. spatial price discrimination. See Section 1.2 for the specifics of the framework, including a review of its relevant variants and the related literature.

Those considerations are then complemented by the second part of the thesis, comprising a single chapter, which considers competition between sellers of a *homogenous* good. A long-standing issue that I address here is the discrepancy between the Cournot model of quantity choice and the alternative price-setting Bertrand specification. A summary of the related literature is provided in Section 1.3.

A central theme of the thesis is the role of uncertainty in addressing the aforementioned controversies. In Part I of the study, this takes the form of uncertainty about the consumers' preferences regarding the characteristics of the heterogeneous products on offer, while in Part II I consider strategic uncertainty about the behaviour of competitors, rather than the potential buyers. See Section 1.4 for an outline of the ensuing discussion and its key results.

1.2 Literature Review: Heterogeneous Goods

In the classic Hotelling [1929] model of spatial competition, consumers are assumed to be uniformly distributed on a unit interval, commonly referred to as the 'Hotelling Main Street'. Although this suggests a strictly spatial context, the model can equally be given an alternative interpretation, where the interval represents a continuum of product characteristics, e.g. 'sweetness of a bar of chocolate', and the distribution of consumer tastes (represented by one's ideal product) is still uniform. See Gabszewicz and Thisse [1992] for more details.

Two firms are said to compete by first simultaneously designing their products / choosing locations in the 'Main Street', before going on to compete in prices in the second stage of the game. Such a timing of moves is justified by the fact that prices can be more easily altered than store locations or product features, so they may be adjusted in reaction to the competitor's first stage location choice once it has been observed.

The manner of the ensuing price competition depends crucially on whether it is the firms or the consumers who are in charge of transport within the 'Main Street'. In the former case, sellers are able to quote different prices, including the cost of delivery, at different locations, so each firm's strategy in the second stage of the game is an accordingly specified price schedule, rather than a single price everyone has to pay.

Naturally, every consumer will then purchase the product from the firm offering the

lowest delivery inclusive price at her location, so that we effectively have an equivalent of Bertrand price competition at every point of the ‘Main Street’. Combined with the standard Hotelling assumption of constant unit production costs, this means that, in equilibrium, the firm closer to any given point attracts the entire local demand by marginally undercutting the competitor’s break-even price (equal to the combined cost of production and transport). A firm’s profit margin is then based on the difference between the transportation costs incurred by the rival and those of its own. Consequently, if a firm tries to maximize the total of its eventual equilibrium profits across all consumer locations, it will itself locate so as to minimize the cost of delivering to the consumers. As a result, the unique Subgame-Perfect Nash Equilibrium (SPNE) of the game has the two firms located at the first and third quartile of the market respectively, which minimizes the total of the transport costs incurred and is socially efficient. The details of this variation of the Hotelling framework, known as ‘spatial price discrimination’, can be found in Lederer and Hurter [1986].

Nevertheless, in reality sellers typically do not have a monopoly over transport services. In fact, in the non-spatial context of consumer tastes this form of discriminatory pricing becomes even more difficult to justify, as it is impossible to restrict an offer to people with particular tastes. For this reason, an alternative ‘mill-pricing’ specification, originally proposed by Hotelling, has retained a prominent place in the literature.

This assumes that it is the consumer who travels to a firm’s factory outlet (‘mill’) in order to make a purchase, thereby incurring the transportation costs of getting there on top of the seller’s chosen price, and hence choosing the outlet where the total of those purchasing costs is smaller. Naturally, the firms are then unable to make prices contingent on the buyer’s location, so their strategy in the second stage of the game is a single ‘mill’ price, rather than a price schedule. An outcome of a particular choice of locations and prices by the duopolists is illustrated in Figure 1.1.

Initially, it was postulated by Hotelling that in order to maximize the resulting second-stage equilibrium profits, the duopolists would want to locate as close as possible to each other - the so called ‘minimum differentiation’ result. This was demonstrated to be incorrect by d’Aspremont, Gabszewicz, and Thisse [1979], who showed that for the classic linear transport cost specification second stage pure-strategy Nash Equilibria (PSE) cease to exist when the firms are located sufficiently close together. However, the authors noted that this is no longer the case for an alternative quadratic cost formulation, which, in addition, leads to a strikingly different, ‘maximum differentiation’ SPNE outcome, where the firms are located at the opposite ends of the ‘Main Street’.

In the traditional ‘spatial’ context, quadratic costs may seem somewhat less appropriate

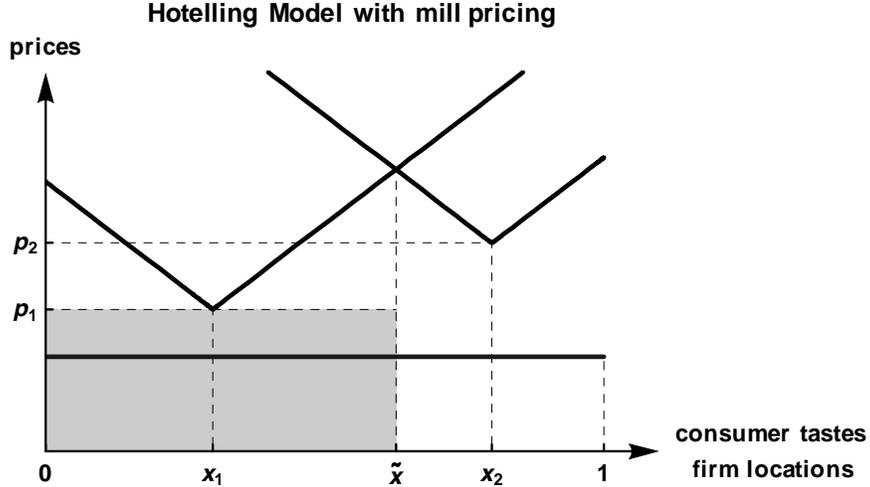


Figure 1.1: Two firms, located at x_1 and x_2 respectively, have set their corresponding prices to p_1 and p_2 . For linear transportation costs, everyone to the left of \tilde{x} finds it cheaper to buy from firm 1, giving it a total revenue of $p_1\tilde{x}$ (shaded region).

than their linear alternative. However, this kind of formulation is more convincing within the more general ‘product characteristics’ scenario, where it reflects a monetary measure of one’s disutility from using a non-ideal product. Indeed, it may well be that compromising on one’s tastes becomes more difficult to bear as one deviates further from the ideal.

The ‘maximum differentiation’ result that follows can also be seen as intuitive. By making their products more different from one another, the firms ensure that they are no longer in direct head-to-head competition at the pricing stage, which allows them to price instead as if they were local monopolists. In fact, Lambertini [1994] showed that the duopolists would go so far as to locate outside the market boundaries when permitted to do so, strategically differentiating their products even when that means making them less attractive to every single consumer.

Nevertheless, the above outcome is still somewhat problematic. On the one hand, the reasoning above provides an economic intuition for the fact that products are *excessively* differentiated relative to the previously discussed socially-optimal situation associated with spatial price discrimination. On the other hand, the *maximum* differentiation result appears too extreme to adequately reflect the economic reality.

This could be attributed to the lack of robustness of the theoretical model to changes in the transport cost specification. In particular, for a family of linear-quadratic costs Gabszewicz and Thisse [1986] demonstrated that whenever the weight associated with the linear component of the cost function is positive, there is a range of locations such that no second-stage PSE exists. The issue was analysed in more detail by Anderson [1988], who

reported two very sound negative results. Firstly, a positive linear component of the cost formulation precludes a pure-strategy SPNE. Secondly, in a game with “sufficiently linear” costs and general mixed price-strategies being allowed, there is no symmetric SPNE where the latter mixed strategies are confined to the non-PSE subgames off the equilibrium path. This is to a large extent mirrored by the results of Economides [1986] for a class of power (rather than linear-quadratic) functions, where the SPNE exists, and exhibits maximum differentiation, only for cost specifications sufficiently close to the quadratic case.

The above problem of PSE non-existence, and the lack of pure-strategy SPNE that it induces, cannot be solved by simply allowing for mixed-strategy pricing. On the one hand, mixed strategy equilibria (MSE) in prices are guaranteed to exist as a direct consequence of Theorem 3 in Dasgupta and Maskin [1986a]. On the other hand, mixed strategy pricing is not only questionable on an empirical level, but also difficult to implement in terms of mathematical tractability even in the simplest linear cost case, as evidenced by the work of Osborne and Pitchik [1987]².

Other features of the Hotelling framework, while often appearing overly simplified compared with the economic reality, turn out to be much less problematic than the quadratic cost formulation. In particular, a seminal contribution of Caplin and Nalebuff [1991] provided general conditions for the existence of a PSE in prices. Building on their results, Anderson, Goeree, and Ramer [1997] characterized the unique equilibrium of the full two-stage game for a class of log-concave consumer density functions, and showed that it exhibits excessive product differentiation.³

Even so, this generalization of the uniform distribution case is only possible for the quadratic cost specification. Similarly, it is possible to extend the model to $n > 2$ players when the costs are quadratic - Brenner [2005], but not when they are linear - Economides [1993]. The same applies to extensions of the standard framework to multi-dimensional product characteristics spaces, such as Tabuchi [1994] or Ansari, Economides, and Steckel [1998].

Overall, it appears that the single most problematic feature of the Hotelling framework is its lack of robustness to changes in the transportation cost specification, and the fact that the only functional form that is technically feasible induces an implausible and socially undesirable maximum differentiation outcome. Two further extensions of the model are relevant to addressing this difficulty in the present thesis.

Firstly, in the standard framework it is assumed that the consumer demand is perfectly

²In fact, allowing for mixed strategies at the location stage is equally complicated and yields an infinite number of equilibria, see Bester, de Palma, Leininger, Thomas, and von Thadden [1996]

³See also Tabuchi and Thisse [1995] for a study of the symmetric triangular consumer distribution and the resulting asymmetric SPNE.

inelastic, because everyone always buys exactly one unit of the good from whichever firm they find cheaper. The existing literature attempts to justify this controversial requirement by arguing that although customers will be unwilling to make a purchase above a certain cost level, this reservation price may be sufficiently high so as not to be strategically significant, in the sense that the equilibrium outcome will be the same whether it is imposed or not. The possibility of a *binding* reservation price was rarely considered, with the exception of Economides [1984], who found that it can reduce the range of locations where the PSE in prices fail to exist, albeit at the cost of multiplicity of PSE in some subgames, making the characterization of SPNE problematic.

Secondly, a relatively recent strand of literature considers the possibility of demand uncertainty, i.e. the consumer preferences being subject to random shocks. One of the very first was a study by Balvers and Szerb [1996], investigating the impact of variations of the products' relative desirability when firms are risk averse and prices fixed. The last assumption means the model is effectively an extension of the location-only variant of the Hotelling framework, as analysed in Eaton and Lipsey [1975]. Hence, it is perhaps more closely related to the literature on electoral competition with probabilistic voting, see e.g. Bernhardt, Duggan, and Squintani [2007]. However, Moscarini and Ottaviani [2001] consider a similar form of uncertainty with respect to the relative quality of the goods, while allowing for price competition.

Another way of introducing uncertainty into the Hotelling framework, one more relevant to the present study, is via the horizontal (rather than vertical, quality related) dimension. In other words, instead of finding a product more/less attractive depending on the outcome of the uncertainty but regardless of what the product is actually like, the consumers change their preference ordering over the range of feasible products as a result of random events. This can be given a more natural economic interpretation, for instance, people may come to like warmer clothes if the weather turns cold, while it is more difficult (though not impossible) to observe a random event that would make them favour the clothes of a particular designer regardless of what they are like.

The first study to have utilized such a specification of demand uncertainty was conducted by Harter [1997], where the uniform customer distribution undergoes an also uniformly distributed random shift. To put it differently, consumer tastes are uniform on an interval of a fixed length, that is equally likely to be placed anywhere inside the space of feasible product characteristics. Firms are assumed to locate sequentially, before the exact distribution of tastes becomes known, at which point simultaneous move price competition ensues. The author reports an increased equilibrium level of product differentiation, compared with a situation in which the outcome of uncertainty is already known at the time of

choosing locations. The result was later reproduced by Casado-Izaga [2000] for the same uncertainty specification, but firms locating simultaneously, as in the classic Hotelling setting. A further generalization of the framework is due to Meagher and Zauner [2005], who parametrized the support of the uniform distribution of the random shock, i.e. the range of feasible tastes in which the unit-length consumer cluster is to be located. In this manner, they provide a continuous version of the previous two authors' comparative statics results, which were based on comparing only two extreme cases. In another study, Meagher and Zauner [2004] extended the results to a random shift that is arbitrarily (rather than uniformly) distributed on a relatively small interval, thereby restricting the model to small degrees of uncertainty.

Finally, an interpretational caveat worth mentioning is that the appropriate measure of product variety in location models is usually the total transportation cost (or disutility) incurred by the consumers due to the mismatch between the products on offer and the consumers' ideal. In settings with endogenized entry decisions, such as Salop [1979] or Anderson, Palma, and Nesterov [1995], product variety can be socially excessive, i.e. firms continue to enter the market even when the cost of doing so exceeds the benefit to the consumers⁴. However, this measure of product variety should not be confused with product differentiation in a Hotelling duopoly. In the latter model the number of firms is fixed, so that it is always socially desirable to have more variety, in the sense that an average consumer can find a product closer to her ideal. Whenever the firms shift from the opposite ends of the market (as in the 'maximum differentiation' outcome) towards the socially optimal positions at the market quartiles, product differentiation decreases but variety improves.

1.3 Literature Review: Homogeneous Goods

The first game-theoretic model of oligopolistic competition, predating game theory itself, was due to Cournot [1838]. Each firm is assumed to independently choose the quantity of a homogeneous good to bring to the market, which is then sold at the demand price corresponding to the firms' aggregate quantity supplied. A Nash Equilibrium is known to exist in this model under very general conditions: it is sufficient that the firms' profits are quasi-concave in own output, or indeed that the marginal revenues with respect to own output are non-increasing in the output of the competitors. See Novshek [1985] for the details. Although the uniqueness of equilibrium is more complicated, not only does it hold

⁴Indeed, with free entry this tendency occurs even under a discriminatory pricing regime, see Bhaskar and To [2004].

for a fairly wide range of demand / cost specifications, but it also displays empirically plausible comparative statics properties with respect to changes in the firms' marginal profitability - Dixit [1986], or an increase in the number of firms, which generally makes the market more competitive - Novshek [1980].

Nevertheless, the Cournot model was often described as being 'right for the wrong reasons' (a term coined by Fellner). Although it yields an outcome consistent with the stylized facts, it does so without allowing for any pricing decisions, whereas in reality firms are at liberty to set both quantities and prices.

In contrast, the model proposed by Bertrand [1883] takes the price to be the dominant strategic variable, where the lowest-priced firm is committed to satisfying the entire corresponding demand. For constant and symmetric marginal costs of production, this results in a well known paradoxical outcome, where all firms price at the cost level, or, for constant but asymmetric costs, the most efficient firm marginally undercuts the break-even price of the second-best one. See Harrington [1989] for a modern treatment, including the subtleties associated with discrete pricing. In fact, the paradox can extend to discontinuous demand / cost specifications, as long as the tie breaking rule (applied when two or more firms charge the same lowest price) allocates the entire demand to one firm ('winner-takes-all'), instead of the usual 'equal-sharing' assumption - Baye and Morgan [2002].

Nevertheless, a direct parallel can be drawn between the 'Bertrand paradox' and the previously discussed 'maximum differentiation' result. Both are empirically questionable, the former since in reality two firms is usually not a number sufficient to guarantee a perfectly competitive outcome, see Bresnahan [1989]. Both are also highly non-robust to changes in the underlying assumptions (in this case, the assumption of constant marginal costs) which result in an indeterminate outcome. For instance, an early critique by Edgeworth [1925] illustrates the non-existence of pure-strategy equilibrium under decreasing returns to scale, particularly in the presence of capacity constraints⁵. When the costs are strictly convex, an opposite problem of PSE multiplicity makes the outcome equally indeterminate Dastidar [1995].

Those difficulties extend to *increasing* returns to scale, where competition drives prices down to the average cost level, which in turn cannot be an equilibrium, since having to share the resulting demand between them firms fail to cover their costs. Indeed, introducing even an arbitrarily small fixed cost to the constant marginal cost model also destroys the PSE, see Shapiro [1989], unless the fixed costs are asymmetric and avoidable when no

⁵Although see Borgers [1992] for a study of how the set of potential outcomes may be reduced via iterated elimination of dominated strategies.

output is sold - Marquez [1997].

As in the Hotelling case, the problem may be avoided by resorting to MSE, which are again guaranteed to exist - Maskin [1986]. Even so, the usual criticism of MSE still applies - Rubinstein [1991], while there is also a possibility of an infinite number of equilibria - Hoernig [2007].

Overall, we therefore have the Cournot model with its mostly determinate, robust and empirically plausible solutions based on a seemingly unrealistic set-up, contrasted with the Bertrand framework, where the strategy set specification is somewhat more convincing, but the outcome either paradoxical or, in most cases, indeterminate.

Unsurprisingly then, considerable research effort has been devoted to resolving the difficulties associated with Cournot. This could entail demonstrating that although in reality firms do not directly compete in quantities alone, their seemingly more complex strategic interaction implicitly complies with the simplified Cournot framework or, as Tirole [1988] put it, has the Cournot reduced form. Alternatively, one could attempt to prove a somewhat weaker claim, that the equilibria of a more general competition framework coincide with those of the Cournot model.

The latter approach is exemplified by Klemperer and Meyer [1986], who allowed the firms to choose either a ‘quantity’ or a ‘price’ strategy. Multiple equilibria follow, involving all possible configurations of strategy types, unless, interestingly, the firms are faced with uncertainty about the market demand. In this case, equilibria based on ‘quantity’ strategies prevail when the marginal costs are relatively steep, which mirrors and extends the earlier results by Bresnahan [1981].

If it is not entirely convincing to allow the firms to choose between different quantity strategies (as in Cournot), it is even less obvious why they might be able to set a quantity *as an alternative* for specifying a price, rather than simply to decide on a quantity-price pair. Fortunately, the above analysis can be equally carried out within a much more general and somewhat more empirically justifiable framework, in which firms choose ‘supply functions’ as their strategies. The motivation for this is that, in reality, producers may commit to a schedule of price-contingent outputs, by means of choosing the firm’s size, structure and setting up the employee incentive system. By restricting the set of available supply functions to ones that never result in a loss, Grossman [1981] was able to show that when the fixed costs are high and the number of firms too large for all of them to achieve a profit, the competitive outcome is the unique equilibrium. However, if any of those conditions is violated, there is an infinite number of equilibria.

The last problem is a consequence of the richness of the players’ strategy spaces and the fact that, for a given strategy profile, there may be an infinite number of unilateral

deviations that would result in the player in question receiving exactly the same payoff. Nevertheless, Klemperer and Meyer [1989] once again demonstrate that the number of equilibria is dramatically reduced in the presence of demand uncertainty. In fact, their results above continue to hold, in the sense that when the marginal cost curves are steeper, then so are the equilibrium supply functions, i.e. they more closely resemble the Cournot commitment to supply a given quantity of output regardless of the price. More recently, Delgado [2006] and Delgado and Moreno [2004] demonstrated how adding pre-play communication to a game with players restricted to non-decreasing supply functions may leave the Cournot outcome as the only coalition-proof equilibrium in the sense of Bernheim, Peleg, and Whinston [1987].

The above cases aside, multiplicity of equilibria is still prevalent in models based on supply functions, making them rather unpopular with theorists. Unfortunately, a simpler and empirically plausible setting in which each firm simultaneously chooses a quantity-price pair exhibits an opposite problem of a lack of a pure-strategy equilibrium - Friedman [1988]. As shown by Dixon [1992] in the context of an Edgeworthian duopoly, the equilibrium non-existence problem can be avoided if the production costs are incurred only for the quantity of output that is eventually sold, rather than the one that is offered for sale as part of the player's strategy choice. Unlike in the Cournot model (where the entire outputs are always sold), the latter quantity may be strictly larger than the former in general price-quantity games. It turns out that when costs are 'sales-dependent' (rather than 'sunk') in the above manner, a unique equilibrium exists and is competitive, as in the classic Bertrand game.

The last result is, of course, hardly a voice in favour of Cournot competition. However, another line of research is more successful in this regard, thanks to introducing dynamics into the original, static problem. In particular, Singh and Vives [1984] consider a two-stage game in which each firm first commits to a price or quantity mode, and only then specifies its exact strategy accordingly. It turns out that choosing the quantity mode is a dominant strategy in the reduced stage one game, although it is not clear how to justify the underlying commitments in reality. A more convincing motivation for the adapted dynamic framework was provided by Kreps and Scheinkman [1983], who study a two-stage game of production capacity choice, followed by price competition subject to the imposed capacity constraints, i.e. in the sense of Edgeworth. Indeed, building production capacity cannot occur instantaneously, as opposed to adjusting the price, justifying the timing of the model via an argument similar to the one discussed for the Hotelling framework. Interestingly, Kreps and Scheinkman show that the Cournot outcome is a SPNE of the game, in the sense that firms choose the equilibrium Cournot quantities as their capacities in the first

stage, and then independently set their prices equal to the demand price associated with the aggregate equilibrium output.

Unfortunately, Davidson and Deneckere [1986] demonstrated that this result is reliant on the so called ‘efficient’ rationing rule, applied when the lower-priced firm is unable to satisfy the entire demand at its price due to the previously self-imposed capacity commitment. In particular, the rule requires that the residual demand of the more expensive firm is minimized and equal to the market demand at its price reduced by the other firm’s capacity. Davidson and Deneckere consider an alternative, ‘proportional’ rationing rule, whereby the higher-priced firm only loses a fraction of the corresponding market demand equal to the proportion of the market demand associated with the lower price satisfied by the cheaper seller. They show that in this case, as well as for most intermediate rationing rules between the ‘proportional’ and ‘efficient’ extremes, the Cournot outcome is no longer sustained as a SPNE.

Nevertheless, the criticism is irrelevant when the demand is perfectly inelastic, as in Acemoglu, Bimpikis, and Ozdaglar [2009] or Fabra, von der Fehr, and Harbord [2006], since when the market demand is the same regardless of the price the two alternative rationing rules coincide. It also does not apply to a uniformly elastic demand, as in Madden [1998], in which case the two-stage game reduces to the Cournot model, given Nash Equilibrium play at the pricing stage and as long as costs are ‘sunk’ in the sense explained above. In other words, *any* profile of first stage quantity / capacity choices results in the second-stage prices being equal to the demand price associated with the aggregate output. Thus, not only the Cournot *outcome* coincides with the equilibrium of a more complex and empirically plausible setting, but the entire Cournot *game* is justifiable as its simplification.

While the last setting may be viewed as a relatively ‘special’ case, a somewhat similar result has more recently been obtained by Moreno and Ubeda [2006], using a model particularly relevant to the present thesis. The framework is essentially the same as in Kreps and Scheinkman, except there are possibly more than two players, who name minimum (rather than exact) prices at which they are willing to sell their corresponding outputs. This is reflected in the way in which the market equilibrium, and hence the firms’ payoffs, are determined. Each firm’s quantity-price pair is effectively an individual supply schedule, offering to provide the chosen quantity if the price is at least equal to the one specified, and zero output otherwise. Technically, this is similar to the ‘supply functions’ framework, except players are restricted to a specific class of stepwise supply correspondences. Consequently, unlike in the Kreps and Scheinkman specification, the entire trade occurs at the market clearing price, given the firms’ *aggregate* supply, where the lower-priced firm



Figure 1.2: Two firms set outputs to q_1 , q_2 , and the respective minimum prices to p_1 , p_2 . The market-clearing price p^* is where the resulting stepwise-increasing supply intersects the downward-sloping demand. Here, the lower-priced firm satisfies the demand at p^* , earning $q_1 p^*$ and leaving rival with nothing.

always has priority to sell. See Figure 1.2 for a two-player example and Section 4.3.1 for a more detailed, formal explanation of the model.

Two features of this framework are particularly important for the current study. First, since all output is sold at the same price, a rationing rule is no longer necessary, eliminating the associated controversies. Second, competition based on minimum (rather than exact) prices can accommodate the complete sets of both the Cournot and Bertrand outcomes. More specifically, when the firms' reservation prices are sufficiently low, their aggregate quantity is sold at the corresponding demand price, exactly as in the Cournot model. In contrast, when the prices are high enough, the least expensive supplier satisfies the entire market demand at her chosen price, just as the Bertrand framework would require.

Moreno and Ubeda improve on the Kreps and Scheinkman results, by demonstrating that, in their model, a pure strategy second-stage equilibrium in prices no longer fails to exist for some capacity choices. Although multiple PSE leading to different outcomes may still arise, the Cournot outcome is the only pure-strategy SPNE of the full two-stage game. Interestingly, the authors point to utility industries as one area where the alternative model specification they put forward is highly relevant. Indeed, recent empirical studies of such markets, such as Puller [2007], report industrial behaviour consistent with the

Cournot framework, suggesting that reservation-pricing may provide a key to a convincing justification and interpretation of the Cournot model.

1.4 Outline of Discussion

I begin the discussion in Chapter 2, by arguing that the most comprehensive Hotelling-based analysis of demand uncertainty to date, conducted by Meagher and Zauner [2005], is still only viable for uncertainty being sufficiently small. This is because uncertainty is modelled via spreading the space of feasible product characteristics / tastes in which the fixed-size consumer cluster is to be located, thereby inducing the firms to increase product differentiation and prices accordingly. Eventually, any finite consumer reservation price is exceeded, thus contradicting the underlying assumption of the said threshold being sufficiently high so as not to be strategically significant.

To address the difficulty, I propose a re-formulation of the model, in which uncertainty is instead modelled via *decreasing* the size of the consumer cluster to be located within a *fixed* space of feasible product characteristics. I also consider a more general linear-quadratic transport cost specification, based on Anderson [1988].

Naturally, when the costs are not fully quadratic, the second-stage PSE in prices still fail to exist in some subgames. However, even when firm locations are sufficiently close, so that there is no PSE for a range of outcomes of the uncertainty, the probability measure of this range decreases as uncertainty becomes higher. As a result, whatever the players hope to earn in the non-PSE subgames (within the limits imposed by the presence of a *finite* consumer reservation price) becomes less important for their expected payoffs. Consequently, for demand uncertainty sufficiently high, I am able to characterize the only SPNE where firms set PSE prices where available, including the equilibrium path. In other words, the equilibrium locations lead to PSE play for any outcome of the uncertainty, while no unilateral deviation is profitable, even when hoping to receive maximum payoffs for any outcomes of the uncertainty where the PSE cease to exist as a result of the deviating player's relocation.

The most interesting aspect of the obtained equilibrium solution is that the associated locations approach the socially-efficient ones at the market quartiles in the limit as uncertainty increases, regardless of which combination of linear and quadratic costs is applied. Thus, the two major discrepancies within the Hotelling framework, between linear and quadratic costs, and mill pricing vs. spatial price discrimination, are asymptotically resolved when uncertainty is large.

Those results are in large part extended in Chapter 3 to a situation in which the firms

do not know the exact distribution of random factors affecting their interaction at the pricing stage. This entails an unknown joint probability distribution of both the demand location (as in Chapter 2) and its price elasticity, represented by the scale of (quadratic) transportation costs. The firms are assumed to resolve the resulting ambiguity using the Arrow and Hurwicz [1972] α -maxmin criterion, i.e. they choose locations so as to maximize the weighted average of the highest and lowest possible second-stage PSE profits, instead of the expected value of those profits with respect to the known probability distribution of the random shock. Thus, the weight associated with the lowest-profit scenario is said to constitute a measure of a firm's pessimism.

It turns out that when the players are sufficiently pessimistic, the presence of demand location uncertainty makes the market more competitive (as it was seen in Chapter 2), in the sense that the unique pure-strategy SPNE exhibits less equilibrium product differentiation and lower prices, regardless of the exact outcome of the random events. However, this effect is dampened by the uncertainty about the elasticity of consumer demand.

After discussing the results and examples of real-world industrial behaviour they might explain, I compare them with the results of Meagher and Zauner, and finally consider two possible extensions of the model. First, a change in the timing of moves, such that both stages of the game are played before uncertainty is resolved, results in a continuum of symmetric PSE in prices when the locations are not too asymmetric. However, in line with the previous results, uncertainty decreases both the maximal and minimal equilibrium prices when the players are pessimistic enough.

The other extension considered is a possibility of a *unilateral* increase of a firm's pessimism. Interestingly, this improves the profits of that firm at the competitor's expense, regardless of the actual resolution of the uncertainty. However, when *both* firms become more pessimistic, they locate closer together and are worse off due to intensified price competition, generating a potential Prisoner's Dilemma situation. Based on this, the Chapter concludes with discussing how the most competitive outcome, in which the duopolists are entirely focused on the worst-case scenario, may be implemented via strategic commitment, elimination of underperforming firms or as a result of corporate taxation.

Chapter 4 and the second part of the thesis begin with a formal definition of the proposed model evaluation criterion for non-cooperative game theory. The concept is based on the idea that players form conjectures about the counterparts' behaviour based on their knowledge of the game's structure, which therefore must be determined prior to the formation of these conjectures. In particular, a restriction of the strategy sets is said to be a stable convention if the excluded strategies are the ones which a rational player would never choose, based exclusively on the belief that others will select *some* strategies

within their restricted choice sets. Given this form of strategic uncertainty, a strategy is irrational when it is weakly dominated by some conventional strategy, as long as others do not behave in an unconventional manner.

After illustrating the above definition on an example, I demonstrate that it constitutes a generalization of the Nash Equilibrium and a special case of the von Neumann - Morgenstern stable set for an accordingly defined dominance relation on the set of strategy profiles. The remainder of Chapter 4 is committed to applying the concept to a simultaneous-move version of the Moreno and Ubeda [2006] model with possibly asymmetric costs.

When the production costs are ‘sunk’, it turns out that a restriction of the players’ quantity-price strategy sets equivalent to a Cournot game is a stable convention if and only if the corresponding Cournot game satisfies two requirements. First, every quantity below the monopoly-optimal one is rationalizable. Second, the marginal revenue of each firm is non-negative even when everyone produces their monopoly-optimal outputs.

In contrast, Cournot competition never constitutes a stable convention when the production costs are ‘sales-dependent’, suggesting that potential difficulties with recovering the costs of unsold output are instrumental in sustaining this mode of oligopolistic interaction. Indeed, for constant (rather than generally non-decreasing) marginal costs, I am able to consider a more realistic mixture of ‘sunk’ and ‘sales-dependent’ costs, showing that the transition from the former to the latter induces a gradual increase in the stringency of the above stability requirements. Having illustrated this with a linear-demand example, I also report that those considerations are irrelevant for the discussion of Bertrand competition, which is never a stable convention.

Part I

Heterogeneous Goods

On the existence and social optimality of equilibria in a Hotelling game with uncertain demand and linear-quadratic costs

The contents of this chapter were published under the same title in

The B.E. Journal of Theoretical Economics: Vol. 11 : Iss. 1

2.1 Introduction

Ever since the classic paper by Hotelling [1929], spatial product differentiation has been a long debated issue in economic literature. The initial “Minimum Differentiation Principle” was overturned by d’Aspremont et al. [1979]. They showed that in a location-then-price duopoly with linear transportation costs and mill-pricing a pure-strategy subgame-perfect equilibrium (SPNE) does not, in fact, exist. This is because the firms’ tendency to agglomerate increases the incentive to price-undercut the rival, destroying the pure-strategy equilibrium (PSE) in prices. The outcome is completely different when the transportation costs are quadratic, i.e. a SPNE exists in which the firms maximize product differentiation.

The problem of PSE non-existence is by no means limited to linear costs. In another study, Gabszewicz and Thisse [1986] observed that, for a family of linear-quadratic costs, there is a range of locations with no second-stage PSE whenever the weight associated with the linear component of the cost function is positive. Anderson [1988] analysed this in more detail, reporting, among others, two very sound negative results. First, a non-zero linear component of the cost specification rules out the existence of a pure-strategy SPNE. Second, in a game with “sufficiently linear” costs and mixed price-strategies allowed, there

can be no symmetric SPNE in which the latter mixed strategies are confined to the non-PSE subgames off the equilibrium path.

Anderson then obtained the candidate equilibrium locations as best-responses to one another among all strategies that permit a second-stage PSE. He conjectured that for costs “not too linear” these locations form the SPNE of the two-stage game, and suggested this could be shown by deriving the mixed-strategy equilibrium payoffs in the non-PSE subgames (or putting some well-behaved upper bound on these payoffs). Unfortunately, as demonstrated by Osborne and Pitchik [1987], this is an extremely difficult task even in the simple linear case. As a result, to this day neither the original Hotelling linear cost specification, nor in fact any cost function with a non-zero degree of linearity, admit of an analytical SPNE solution.

These problems no longer appear with mill-pricing replaced by spatial price discrimination, see Gabszewicz and Thisse [1992] or Lederer and Hurter [1986]. Here, the socially-optimal locations (at the market quartiles) obtain in equilibrium, regardless of which combination of linear and quadratic costs is applied. However, this form of price discrimination is not always available and is difficult to justify when the firms compete in a space of consumer tastes. The purpose of this paper is to argue that in the absence of spatial price discrimination a substantial enough degree of demand uncertainty could play a similar role in ensuring social optimality and robustness to altering the usual quadratic transport cost specification. In other words, the much-debated contrasts between the two main Hotelling pricing schemes, as well as the two most common transportation cost functions (linear vs. quadratic), vanish for demand uncertainty sufficiently high.

Several papers have already introduced some form of demand uncertainty into a modified Hotelling setting. Balvers and Szerb [1996] study the effect of random shocks to the products’ desirability under fixed prices. Harter [1997] examines the uncertainty in the form of a uniformly distributed random shift of the (uniform) customer distribution, where the firms locate sequentially. Other papers, such as Aghion, Espinosa, and Jullien [1993], concentrate on the strategic effect of acquiring information about the demand through price-experimentation.

On the other hand, relatively few studies consider the effect of demand uncertainty in an otherwise unchanged Hotelling framework. Of those, Casado-Izaga [2000] adopts the same form of uncertainty as Harter, but the duopolists locate simultaneously, prior to observing the actual customer distribution and then choosing prices. Meagher and Zauner [2005] consider a similar setting, but succeed in parameterizing the support of the (uniform) random variable that shifts the customer distribution and report that demand uncertainty increases the equilibrium level of product differentiation. The current framework is almost

the same, with the exception of a more general cost specification and some subtle changes, essential in establishing the analogy with price discrimination and detailed in section two of the paper.

It would appear that uncertainty could only intensify the problems (as discussed above) associated with the mill-pricing regime. It is more difficult to ensure a PSE not only for various location-pairs, but also for all realizations of customer demand. However, uncertainty makes it possible to follow Anderson’s suggestion and put a tractable upper bound on the mixed strategy equilibrium (MSE) payoffs, in the form of the optimal monopoly profit (subject to the imposed consumer reservation price). This would be insufficient under certainty, but in the present context the non-PSE subgames are becoming more scarce when there is more uncertainty, meaning that the associated second-stage profits are less and less important for the first-stage expected payoff. Consequently, even a “generous” upper bound is eventually sufficient to show that no deviation from the candidate equilibrium locations is profitable.

As a result, a symmetric closed-form solution is obtained, based on the linear-quadratic cost function. For demand uncertainty sufficiently high, it is shown to constitute a SPNE as postulated by Anderson, i.e. with mixed price-strategies played only in the non-PSE subgames off the equilibrium path. However, Anderson’s second negative result does not extend here, since the equilibrium exists for any “degree of linearity” of the cost function. Furthermore, as argued in section three, the current result also has an exclusively pure-strategy interpretation.

The obtained equilibrium solution is particularly interesting, because as the demand becomes more uncertain, the equilibrium locations of the uncertainty mill-pricing game converge to the social-optimum associated with the SPNE of the certainty price discrimination model. In fact, the discriminatory pricing game is shown to constitute a limiting case of the current model under the most general conditions. In other words, demand uncertainty (asymptotically) reconciles the inconsistent outcomes of linear and quadratic costs under mill-pricing with the socially-efficient outcome of discriminatory pricing.

2.2 The Model

As in the classic framework by Hotelling [1929], there are two firms simultaneously choosing locations $x_1, x_2 \in [0, 1]$ in the first stage of the game and prices p_1, p_2 in the second stage, with a consumer located at x minimizing the total cost of purchase $p_i + c(|x_i - x|)$ over $i = 1, 2$, so long as it does not exceed a finite reservation price r . The transportation cost

function $c(\cdot)$ is linear-quadratic, i.e.:

$$c(|x_i - x|) = a|x_i - x| + (1 - a)|x_i - x|^2, \text{ where } a \in [0, 1] \quad (2.1)$$

As usual, it is assumed that $x_1 \leq x_2$ and the marginal production cost is zero. Furthermore, in the second stage of the game the firms are allowed to employ mixed price-strategies.

Uncertainty is introduced by assuming that the second-stage (ex-post) customer distribution is uniform on an interval $[z, z + m] \subset [0, 1]$, i.e. contained in the space of feasible product characteristics. At the stage of choosing locations, the firms are uncertain of the exact value of z , which they know to be uniformly distributed on the interval $[0, 1 - m]$, where $m \in [0, 1]$ is a parameter. In other words, the duopolists expect the demand to be uniform on a segment of length m , equally likely to be placed anywhere inside the space of feasible product characteristics.

It is clear that $m = 1$ corresponds to the standard “certainty” case. As m becomes smaller, the demand begins to vary between the states of nature. Initially the uncertainty is small, for instance, when $m = 3/4$, the support of the customer distribution is equally likely to be anything from $[0, 3/4]$ to $[1/4, 1]$. However, as m further decreases, the consumer tastes are gradually becoming more variable, until at $m = 0$ they are completely state-specific; everyone is located at the same point, the distribution of which is uniform on $[0, 1]$. Hence, m may be thought of not only as the ex-post differentiation of tastes, but also as “the degree of certainty”.

As mentioned earlier, the uncertainty is resolved between the two stages of the game. The usual solution to the location problem is then to consider a reduced game in which the payoffs associated with any $\{x_1, x_2\}$ are the expected values (with respect to the distribution of z) of the second-stage equilibrium profits associated with this location-pair.

Remark. The above specification of demand uncertainty also has a natural economic interpretation. To see this, suppose the distribution of consumer preferences is initially uniform on $[0, 1]$. The customers then observe a signal indicating exactly what type of product is best and the firms think this suggestion is equally likely to be anything within the space of feasible product characteristics. Having observed a particular value s of the signal, a consumer located at x re-locates to $x + (s - x)(1 - m)$. In other words, everyone moves towards the value of the signal, by a fraction $(1 - m) \in [0, 1]$ of the distance between the value of the signal and their original location. Consequently, people whose initial views were further from what is now suggested as best shift their preferences more than those individuals who were already close to the value of the signal.

It follows that $(1 - m)$ can be interpreted as the “strength of the signal” and it is easy to check that the resulting ex-post customer distribution is uniform on an interval $[(1 - m)s, (1 - m)s + m]$. Since s is uniform on $[0, 1]$, the latter is a segment of length m , equally likely to be placed anywhere inside the space of feasible product characteristics (just as in the original model specification above).

The current approach can be seen therefore as taking the classic Hotelling framework and introducing a random shock/signal which distorts the original preferences in such a way that they are still bound to remain within the initial “main street” of $[0, 1]$. When $m = 1$, the signal’s strength is zero and it has no effect at all. However, as m decreases, the random signal begins to affect the customer preferences more and more, making them more variable and state-specific. A demand more sensitive to random factors means there is more demand uncertainty.

Relationship to existing models. The current “uniform-uniform” specification of demand uncertainty is similar to the one in Meagher and Zauner [2005]. The general framework from which both originate has the consumers uniformly distributed on a segment of length m , shifted by an also uniformly distributed shock with a value of up to L . Consequently, the support of the ex-post customer distribution can be anything from $[0, m]$ to $[L, L + m]$ and the space of feasible tastes (that could appear in some state of nature) is $[0, L + m]$. Using the classic ice-cream sellers example, this entails people uniformly distributed on a segment of length m , equally likely to be found anywhere on a beach of length $L + m$.

In fact, what determines the outcome of the model (subject to an appropriate re-scaling) is the ratio $(L + m)/m$ of the length of the space of feasible tastes to the ex-post length of the market. This is also the most appropriate, invariant to scale measure of demand uncertainty in this context. It is therefore appropriate to reduce the framework to a one-parameter model, which could be done in various ways.

If we model demand uncertainty as an increase of L while holding m fixed (as in Meagher and Zauner [2005]), then more uncertainty means “spreading the beach”, with an unchanged consumer cluster equally likely to be found anywhere inside its new boundaries. Of course, one could simply define $(-\infty, +\infty)$ as the space of feasible products, but another issue is that, as uncertainty increases, the SPNE product differentiation and average cost of purchase increase without a bound, approaching infinity in the limit. Consequently, any finite consumer reservation price, no matter how high, would have to be at some point exceeded. Such a specification would therefore seem appropriate to model demand uncertainty when it is “small”, precisely the opposite of the present paper’s focus of

interest¹.

In contrast, the idea proposed here is to make the demand more or less variable while still “within the same beach”. This is achieved by setting $L = 1 - m$, so as to fix the space of feasible tastes at the traditional unit interval. Consequently, the “uncertainty ratio” is now equal to $1/m$. Instead of “expanding the beach” to make the demand more uncertain, we keep it fixed, while reducing the length of the consumer cluster. In essence, the location of a needle in a haystack is more uncertain than that of a substantially larger object

Constraining the products/tastes to a finite Hotelling “Main Street” also has the advantage of allowing for a consumer reservation price that, despite being finite, is sufficiently high so as not to be strategically important anyway (exactly as it is assumed in the classic Hotelling model). This is not only realistic in economic applications, but will also prove crucial in establishing the equivalence with discriminatory pricing.

2.3 Results

The reasoning behind the paper’s main result below may be outlined as follows. For a convex transportation cost function, a second-stage PSE exists when both firms are located in the exterior of the ex-post customer distribution, i.e. when $x_1, x_2 \notin [z, z + m]$. This is because profits are then concave in own price. On the other hand, if at least one firm located inside the market and the rival located close enough relative to the market’s length, a discontinuity of marginal profits may occur where one firm encroaches on the other’s hinterland. This creates an incentive to undercut the competitor’s price at the candidate equilibrium (one satisfying the first-order conditions for profit maximizing), thus destroying the PSE. See Anderson [1988] for the details.

Nevertheless, if the firms are distant enough from one another relative to m , a second-stage price equilibrium will exist for any $z \in [0, 1 - m]$. This is because for larger product differentiation a lower price is required to advance into the rival’s hinterland and undercutting becomes less attractive.

Hence, the idea is first to identify the “candidate” (local) equilibrium locations, i.e. ones which are best-responses to one another among all locations that permit a pure-

¹Another study by Meagher and Zauner [2004] considers a random shock arbitrarily distributed on a fixed interval. Although product differentiation is no longer unbounded, the model requires the variance of the shock to be small relative to the ex-post differentiation of tastes, so that no firm would ever capture the entire market in *any* state of nature. Consequently, this model is also valid for demand uncertainty not too big.

strategy price equilibrium for all values of z . As m decreases, in order to get close enough to the rival to start destroying the PSE in some subgames, one has to deviate further from the “candidate” locations, losing more expected profits due to increased price competition. Furthermore, even once non-PSE subgames begin to appear, they are less likely to occur for smaller values of m , because the chances of some firm ending up in the interior of a reduced market are lower. For small enough m , even when hoping to earn optimal monopoly profits in all non-PSE subgames, one will be better off at the “candidate” equilibrium position. This is formalized in the following.

Proposition 1 *Consider a two-stage Hotelling duopoly with linear-quadratic transport costs (2.1), where $a \in [0, 1)$, and demand uniform on a segment of length $m \in (0, 1)$, equally likely to be found anywhere in $[0, 1]$. For any non-binding finite reservation price r and a small enough m there exists a subgame-perfect Nash equilibrium in which the unique pure-strategy equilibrium prices are played in all subgames in which they exist, including the equilibrium path, while mixed strategy equilibrium prices are played in all the other subgames. In any such equilibrium, locations are given by:*

$$x_1^* = 1 - x_2^* = \frac{1 + a(1 + 4m) - 8(1 - a)m^2}{4[1 + a - 2(1 - a)m]} \quad (2.2)$$

Proof. See the Appendix. ■

It is interesting to view these findings in the context of the “certainty” mill-pricing model. In the linear case, originally studied by Hotelling, the best-response dynamics takes the duopolists towards the centre of the market, which led him to formulating the “Minimum Differentiation Principle”. This was overturned later in d’Aspremont et al. [1979], by observing that this tendency to reduce product differentiation eventually makes the firms enter the area where no PSE exists.

In contrast, in the linear-quadratic model by Anderson [1988], corresponding to the case of $m = 1$ in the current specification, one can identify candidate equilibrium locations for “not too linear” transportation costs². Those locations are “local” equilibria, in the sense that a firm’s second-stage PSE profit will decline as it moves away from the candidate equilibrium location and towards the rival. However, one can still locate close enough to the competitor so as to eliminate the PSE, making it impossible to show that the proposed locations form a SPNE without obtaining some estimate of the payoffs in the non-PSE subgames.

²Although Anderson uses two parameters, a and b as the respective weights of the linear and quadratic components, it turns out that the results only depend on a/b , i.e. the model is compatible with the current specification.

Similarly, in case of $m \in (0, 1)$ studied in Proposition 1 one can identify “local” equilibrium locations such that locating any closer to the counterpart will decrease the second-stage *expected* profit, as long as PSE in prices exist *for all realizations of the uncertainty* $z \in [0, 1 - m]$. This can be done for $1 - a$ sufficiently large relative to m (“not too linear” costs) or, conversely, for m small enough relative to $1 - a$ ($m \leq 1 / \lceil 2\sqrt{2} \rceil$ also ensures that $x_1^* \in [0, 1/2]$).

In fact, for m sufficiently close to 0, a symmetric SPNE is no longer impossible for “linear enough” costs. The candidate equilibria exist for all $a \in [0, 1)$ and can be shown to form SPNE. In other words, not only does Anderson’s negative result not extend here, but the SPNE with locations x_i^* may be established.

The finite reservation price r is instrumental in proving Proposition 1. In fact, the possibility of consumers choosing an “outside option” at some point has already been studied as a potential remedy for the problem of PSE non-existence under certainty. See, for example, Economides [1984] or, for an extensive recent discussion of the resulting second-stage equilibria, consult Merel and Sexton [2010].

The difference is that the current model follows the mainstream of spatial competition literature in assuming that r is “non-binding”, i.e. sufficiently high so as not to affect the equilibrium characteristics in any way. Despite this, it turns out it can still facilitate establishing the SPNE. This is because the reservation price is used not as a factor affecting the players’ PSE play, but rather as a basis for putting an upper bound on the MSE payoffs when the PSE fail to exist.

Such a “generous” upper bound would still be insufficient in case of $m = 1$, as studied by Anderson. However, the non-PSE subgames become more scarce as m decreases. Hence, any finite benefits of deviating from x_i^* achieved in those subgames are eventually outweighed by the losses due to increased price competition incurred in those subgames in which the PSE still exist.

We now turn to the case of $a = 1$, i.e. the standard linear cost function. This is the case most-exposed to the problem of PSE non-existence. Firstly, the incentives to price-undercut the rival at her location are maximized, since this results in instantaneously acquiring the entire market. Secondly, the cost function is least convex among the studied class, minimizing the tendency for strategic product differentiation, and hence making it easier to undercut the closely located competitor. For more on the relationship between the curvature of transport costs and PSE existence under certainty, see Economides [1986].

The reason why $a = 1$ is a case which requires a separate treatment is that the second-stage profit functions are no longer continuous. However, the proof of Proposition 2 below is analogous to that of Proposition 1 above and the latter result can be extended here. In

fact, because this case is also relatively simple in terms of mathematical tractability, more can be said about the degree of uncertainty that is sufficient for the proposed SPNE to occur.

Proposition 2 *Consider a two-stage Hotelling duopoly with linear transport costs (2.1), where $a = 1$, and demand uniform on a segment of length $m \in (0, 1)$, equally likely to be found anywhere in $[0, 1]$. For any non-binding finite reservation price r and a small enough m there exists a subgame-perfect Nash equilibrium in which the unique pure-strategy equilibrium prices are played in all subgames in which they exist, including the equilibrium path, while mixed strategy equilibrium prices are played in all the other subgames. In any such equilibrium, locations are given by x_i^* with $a = 1$. Furthermore it is sufficient that $m \leq 1/9$ for such an equilibrium to appear.*

Proof. See the Appendix. ■

The above threshold value of m is relatively low. However, there are good reasons to believe that, in general, less uncertainty is required to ensure the existence of the proposed SPNE.

First of all, in establishing the above sufficient condition no attempt was made to obtain the exact MSE payoffs in the non-PSE subgames. Instead, a “generous” upper bound was put on those payoffs, equal to the maximum of what a firm could earn over all price-pairs. This was partly with mathematical tractability in mind and partly to preserve an exclusively pure-strategy interpretation of the obtained SPNE³. Nevertheless, there is evidence to suggest that the MSE payoffs are lower than those in the candidate PSE (satisfying the first-order conditions) and diminishing quickly as product differentiation falls (see Osborne and Pitchik [1987]). In any case, they must be significantly lower than the applied upper bound and hence less uncertainty is required to render a deviation from x_i^* into the non-PSE area unprofitable.

Another reason why, in general, less uncertainty should be necessary for the proposed SPNE to exist, is that, as already indicated, $a = 1$ is the case most-exposed to the problem of PSE non-existence. Consequently, the threshold value of m is likely to be higher than $1/9$ for $a < 1$, despite being more difficult to obtain⁴.

Finally, one could consider binding (i.e. arbitrarily low) reservation prices. Although complicated on the technical side, this would allow for lower upper bounds on the MSE

³By construction, locations x_i^* are best-responses to one another regardless of what happens in any non-PSE subgames. Hence, in a game with pure-strategy pricing, a player i who considers the unique PSE a credible indication of second-stage profits will also regard x_i^* as an optimal reaction to x_{-i}^* .

⁴Indeed, this intuition can be confirmed by substituting any specific value for a in the online Appendix proof of Proposition 1, so as to obtain a numerical approximation of the threshold, despite a generic closed form solution being unavailable.

payoffs, while also potentially reducing the range of locations / states of nature where the PSE fail (see Economides [1984]).

A closer look at the SPNE locations (2.2) reveals that $\partial x_1^*/\partial a > 0$. Not surprisingly, a higher curvature of the cost function results in an increased tendency for strategic product differentiation, just as it happens under certainty. However, we also have $\partial^2 x_1^*/\partial a \partial m > 0$, meaning that this effect is weakened when demand uncertainty increases. In other words, the model is more robust to variations in the transport cost function when the demand becomes more uncertain.

In fact, as $m \rightarrow 0$, i.e. approaches the extreme case all customers located at the same point, we have convergence to $x_1^* = 1 - x_2^* = 1/4$ for any $a \in [0, 1]$ ⁵. This leads to an interesting analogy with the aforementioned 'certainty' model of discriminatory (rather than mill-) pricing.

Corollary 3 *As $m \rightarrow 0$, the equilibrium locations x_i^* converge to the socially optimal ones associated with the certainty price-discrimination model for any $a \in [0, 1]$.*

The next proposition shows that this is by no means a coincidence. In the extreme case of $m = 0$, a direct equivalence between demand uncertainty and discriminatory pricing holds under the most general conditions, in terms of the distribution of the random shock, the transport cost function and the number of players, which can be greater than two⁶.

Proposition 4 *Consider a two-stage Hotelling mill-pricing n -player game with any non-decreasing transport cost function $c(|x_i - x|)$, in which all consumers are located at the same point drawn from a probability distribution F . This game is equivalent to a "certainty" price discrimination Hotelling n -player game with transport cost function $c(|x_i - x|)$ and customer distribution F .*

Proof. Take a second-stage price subgame of the "uncertainty" game in which the consumers are located at point z . This is clearly the same as a Bertrand game with firm i producing at a (asymmetric) cost $c(|x_i - x|)$. Consequently, the firm closest to z (and hence most cost-efficient) captures the entire demand, by marginally undercutting a zero

⁵A similar observation has been made by Kieron Meagher in the context of a quadratic cost game ($a = 0$) with uncertainty specified as in Meagher and Zauner [2005]. This can be found in an unpublished working paper "On the Equivalence of Asymptotic Demand Location Uncertainty and Spatial Price Discrimination". I am grateful to Kieron Meagher for making it available to me.

⁶This refers to the usual n -player Hotelling framework (as in Brenner [2005]). Each consumer minimizes the total purchasing cost $p_i + c(|x_i - x|)$ over $i \in 1, 2, \dots, n$ and the firms simultaneously choose first locations and then prices. Everything else is as defined in Section 2.

mill price of the second-closest firm. The profit of firm i is therefore:

$$\pi_i^*(x_1, \dots, x_n, z) = \max \left\{ 0, \min_{j \neq i} [c(|x_j - z|) - c(|x_i - z|)] \right\}$$

which is the same as the second-stage equilibrium profit attained at location z in the “certainty” price discrimination game (see Gabszewicz and Thisse [1992] or Lederer and Hurter [1986]). Consequently, the first stage expected profit in the “uncertainty” game becomes:

$$\Pi_i(x_1, \dots, x_n) = \int_{-\infty}^{+\infty} \pi_i^*(x_1, \dots, x_n, z) F(z)$$

which, again, is the same as the total profit from all locations in the certainty price discrimination game. As the payoff functions and strategy spaces are the same, the two games are equivalent. ■

The above result provides a helpful insight into the mechanics of the model and makes it possible to interpret the first two propositions in terms of the analogy between demand uncertainty and discriminatory pricing.

When the demand varies between the states of nature, the firms are effectively price discriminating between them, adapting different price strategies for different realizations of the uncertainty. As m decreases, it becomes more likely for the ex-post consumer demand to be located in a firm’s hinterland, while also far enough from the rival for the firm to favour capturing it entirely. Such “monopolistic” equilibria gradually replace the “competitive” ones and the firms’ profits are to an increasing extent stemming from the advantage in transportation costs, i.e. from being better placed relative to the ex-post customer demand. This creates an incentive for the duopolists to minimize the transport costs incurred by the consumers over all states of nature, which means implicitly pursuing a socially-optimal objective. For this reason, the curvature of the cost function also becomes less and less important for location decisions. As m goes to zero, the equilibrium locations x_i^* associated with different values of a converge to the efficient ones associated with the “certainty” price discrimination model.

On the one hand, these results are in contrast with those of Meagher and Zauner. Instead of an increase of product differentiation and a decline of welfare as uncertainty increases, we have convergence to the intermediate, socially optimal differentiation level as uncertainty approaches its upper bound. This is because the current model specification makes it possible for the demand to become more uncertain while always remaining within the same boundaries, so as not to confound the effect of the consumers being on average spread over a larger area and the apparently opposite impact of the uncertainty increase.

On the other hand, it is interesting to note that for large levels of demand uncertainty locations x_i^* are not far from the ones obtained by Osborne and Pitchik [1987] who focused on the mixed-strategy price equilibria under certainty and linear costs. In a sense, demand uncertainty is similar to price-randomization, in that any strategy-pair produces a random distribution of second-stage profit allocations.

In fact, an even more direct parallel can be drawn between the current model and the work of Bester et al. [1996]. The latter study considers mixed-strategy locations under quadratic costs and the resulting PSE in prices. When a player responds to a random distribution of the competitor's locations for a fixed distribution of tastes, it is similar to responding to a fixed location of the competitor when the exact placement of the customer demand is random. Either way, a firm is uncertain of how close it would be to the consumer distribution when compared with the rival, which is what drives the second-stage equilibrium profits. However, in the present model a firm cannot itself set a random response, but is restricted to choosing a single location. It is therefore to no surprise that the former, richer strategy space results in an infinite number of possible equilibrium configurations (see Bester et al. [1996] for the details). What the current approach and the other two mentioned have in common is that they all allow for an intermediate level of product differentiation, possibly more realistic than the extreme maximum or minimum differentiation results.

2.4 Concluding Remarks

The paper examined the Hotelling two-stage mill-pricing duopoly with customer demand uniform on a segment of a given length, equally likely to be found anywhere in the usual unit interval of feasible product characteristics. Parametrizing the segment's length to reflect the degree of uncertainty made it possible to avoid assuming an infinite customer reservation price. This in turn allowed for an investigation of the effect of an arbitrarily uncertain demand.

In particular, based on a linear-quadratic cost function, SPNE were established for demand uncertainty sufficiently large, similar to the candidate equilibria postulated by Anderson under certainty. However, Anderson's negative results do not extend here, as the proposed SPNE exist for any degree of linearity of the cost function.

In fact, as demand uncertainty increases, the curvature of transport costs becomes less important for the equilibrium outcome and the SPNE locations associated with different linear-quadratic cost functions converge to the socially-optimal ones associated with the certainty price-discrimination model. In other words, the much-debated contrasts between

the two main Hotelling pricing schemes, as well as the two most common transportation cost functions (linear vs. quadratic), vanish for demand uncertainty sufficiently high.

The intuitive explanation of those results relies upon observing that, when setting different prices for different realizations of customer demand, the duopolists are effectively price-discriminating between the states of nature. When the ex-post differentiation of tastes decreases, it becomes vital to secure the equivalent of a cost advantage over the opponent, by making it cheaper for the customers to get to the firm’s location. Eventually, the flexibility of price-discriminating between the states of nature becomes close to that of setting prices independently at each physical location and the firms implicitly pursue the socially desirable objective of minimizing the total transportation costs. In the limiting case of consumer tastes being completely state-specific, the “uncertainty” mill-pricing model and the “certainty” discriminatory pricing game are equivalent under the most general conditions.

Appendix

Some algebraic derivations are relegated to the on-line Wolfram Mathematica appendix (<http://tinyurl.com/29vk5v1>), where numbers in curly brackets {#} represent a position therein where a particular statement is verified. A free Wolfram Mathematica Reader software is available for download at: <http://www.wolfram.com/products/player/>, and a non-interactive PDF version of the on-line appendix is available at: <http://tinyurl.com/4x9o5du>.

Proof of Proposition 1. We begin by obtaining the “candidate” equilibrium locations (2.2). To this end, consider first a second-stage subgame associated with a particular location-pair and value z of the random shock. Assume, for the moment, that the firms are restricted to playing pure-strategy prices and let \tilde{x}_i denote the value of x solving equation (A*i*) below:

$$p_1 + a(x_1 - x) + (1 - a)(x_1 - x)^2 = p_2 + a(x_2 - x) + (1 - a)(x_2 - x)^2 \quad (\text{A1})$$

$$p_1 + a(x - x_1) + (1 - a)(x_1 - x)^2 = p_2 + a(x - x_2) + (1 - a)(x_2 - x)^2 \quad (\text{A2})$$

$$p_1 + a(x - x_1) + (1 - a)(x_1 - x)^2 = p_2 + a(x_2 - x) + (1 - a)(x_2 - x)^2 \quad (\text{A3})$$

where (A1) and (A2) correspond to the hinterlands of the respective firms and (A3) to the area “in between”. Next, define $p_{i,j}^c$ as the candidate “competitive” equilibrium price

of firm i in area j , i.e. as the solution to:

$$\frac{\partial [p_1 (\tilde{x}_j - z) / m]}{\partial p_1} (p_{1,j}^c, p_{2,j}^c) = \frac{\partial \{p_2 [1 - (\tilde{x}_j - z) / m]\}}{\partial p_2} (p_{1,j}^c, p_{2,j}^c) = 0$$

where $\pi_{i,j}^c$ is the corresponding candidate equilibrium profit. Similarly, let $p_{i,j}^m$ denote the candidate “monopolistic” equilibrium price of firm i in area j , i.e. the solution to $\tilde{x}_j (p_{1,j}^m, 0) = z + m$ for player 1 and $\tilde{x}_j (0, p_{2,j}^m) = z$ for player 2. For “exterior” locations, i.e. $x_1, x_2 \notin [z, z + m]$, the profit functions are concave and a PSE must exist. Indeed, in case of $z + m \leq x_1$ we have {1}:

$$\frac{\partial [p_1 (\tilde{x}_1 - z) / m]}{\partial p_1} (p_{1,1}^m, 0) < 0 \Leftrightarrow z < z_{1,1} = \frac{1}{2} \left(\frac{a}{1-a} + x_1 + x_2 \right) - 2m$$

i.e. for $z < z_{1,1}$ firm 1 chooses to take the whole market even when $p_2 = 0$, earning $p_{1,1}^m$. On the other hand, for $z \geq z_1^1$ a competitive equilibrium is “feasible”, i.e.:

$$\tilde{x}_1 (p_{1,1}^c, p_{2,1}^c) \leq z + m \{2\}$$

By symmetry, for $z > x_2$ and $z > z_{2,2} = (x_1 + x_2 - a / [1 - a]) / 2 + m$ firm 2 monopolizes the market, while for $x_2 < z < z_{2,2}$ profits are given by $\pi_{i,2}^c$.

When the market is not contained in some firm’s hinterland, we have {3}:

$$\begin{aligned} \frac{\partial p_1 [p_1 (\tilde{x}_3 - z) / m]}{\partial p_1} (p_{1,3}^m, 0) < 0 &\Leftrightarrow z < z_{1,3} = \frac{1}{2} (x_1 + x_2) - 2m \\ \frac{\partial p_2 [1 - (\tilde{x}_3 - z) / m]}{\partial p_2} (0, p_{2,3}^m) < 0 &\Leftrightarrow z > z_{2,3} = \frac{1}{2} (x_1 + x_2) + m \end{aligned}$$

i.e. for $x_1 - m < z < z_{1,3}$ firm 1 captures the entire market at price $p_{1,3}^m$, while firm 2 does the same for $z_{2,3} < z < x_2$ at price $p_{2,3}^m$. On the other hand, when the firms are placed on the market’s opposite sides (i.e. $x_1 < z$ and $z + m < x_2$) then profits are concave and given by $\pi_{i,3}^c$ in equilibrium for $z_{1,3} < z < z_{2,3}$. Furthermore, for $x_2 - x_1 > 4m$ we have $z_{1,3} > x_1$ and $z_{2,3} < x_2 - m$, which means a PSE exists for any value of z . Suppose that, in addition, $x_1 > m$ and $x_2 < 1 - m$. Then the expected PSE profit of firm 1 is:

$$\Pi_1(x_1, x_2) = \int_0^{x_1-m} p_{1,1}^m dz + \int_{x_1-m}^{z_{1,3}} p_{1,3}^m dz + \int_{z_{1,3}}^{z_{2,3}} \pi_{1,3}^c dz$$

where we have used the fact that $z_{1,3} < z_{1,1}$ and skipped multiplying by $1/(1-m)$, the density of the distribution of z , since it will not affect the optimal choice of location. By

symmetry, we have $\Pi_2(x_1, x_2) = \Pi_1(1 - x_2, 1 - x_1)$ and {4}:

$$\frac{\partial^2 \Pi_1}{\partial x_1^2} < 0, \frac{\partial^2 \Pi_2}{\partial x_2^2} < 0$$

leading to the first order conditions:

$$\frac{\partial \Pi_1}{\partial x_1}(x_1^*, x_2^*) = \frac{\partial \Pi_2}{\partial x_2}(x_1^*, x_2^*) = 0$$

The only solution to those FOC's which could satisfy

$$x_2^* - x_1^* > 4m \text{ and } m < x_1^* < x_2^* < 1 - m$$

is (2.2), and a sufficient condition for this to happen is $m \leq 1/10$ {5}.

Let Π^* denote the expected profit obtained by substituting (2.2) into Π_1 . We will show that for m small enough it is not profitable for player 1 to deviate from x_1^* , regardless of what the players do in the non-PSE subgames.

Consider first a deviation to $x_1 < m$, in which case the expected profit is:

$$\widehat{\Pi}_1^1(x_1, x_2^*) = \int_0^{z_{1,3}} p_{1,3}^m dz + \int_{z_{1,3}}^{z_{2,3}} \pi_{1,3}^c dz$$

and we have {6}:

$$\frac{\partial \widehat{\Pi}_1^1}{\partial x_1}(x_1, x_2^*) > 0 \text{ for } x_1 < m < 1/10$$

Next, consider the case of $x_2^* - 4m < x_1 < x_2^* - 2m$, so that $x_1 - m < z_{1,3} < x_1$ and $x_2^* - m < z_{2,3} < x_2^*$. Since the profits in any subgame (including the non-PSE ones) cannot exceed r , the overall expected profit must then be smaller than:

$$\int_0^{x_1-m} p_{1,1}^m dz + \int_{x_1-m}^{z_{1,3}} p_{1,3}^m dz + \int_{z_{1,3}}^{x_1} r dz + \int_{x_1}^{x_2^*-m} \pi_{1,3}^c dz + \int_{x_2^*-m}^{z_{2,3}} r dz$$

which in turn is less than:

$$\widehat{\Pi}_1^2(x_1, x_2^*) = \int_0^{x_1-m} p_{1,1}^m dz + \int_{x_1-m}^{x_1} r dz + \int_{z_{1,3}}^{z_{2,3}} \pi_{1,3}^c dz + \int_{x_2^*-m}^{x_2^*} r dz$$

where the fact that $\pi_{1.3}^c \geq 0$ {7} was used. For $m \leq 1/10$, $\widehat{\Pi}_1^2(x_1, x_2^*)$ is concave in x_1 {8}. Solving:

$$\frac{\partial \widehat{\Pi}_1^2}{\partial x_1}(x_1, x_2^*) = 0$$

for x_1 , substituting the obtained value back into $\widehat{\Pi}_1^2(x_1, x_2^*)$, subtracting the outcome from Π^* and finally calculating the limit of the resulting expression as $m \rightarrow 0$ yields $(5 + 7a)/256 > 0$ {9}. This means that for m sufficiently small $\widehat{\Pi}_1^2(x_1, x_2^*)$ is less than Π^* and a deviation from x_1^* to $x_1 \in [x_2^* - 4m, x_2^* - 2m]$ is not profitable.

Next, note that for $m \in [0, a/[2(1-a)]]$ we have $z_{1.1} > x_1 - m$ and $z_{2.2} < x_2^*$ {10} and consider the following cases:

1. {11} $x_1 \in [x_2^* - 2m, x_2^* - m]$:

$$\widehat{\Pi}_1^3(x_1, x_2^*) = \int_0^{x_1-m} p_{1.1}^m dz + \int_{x_1-m}^{x_1} r dz + \int_{x_1}^{x_2^*-m} \pi_{1.3}^c dz + \int_{x_2^*-m}^{x_2^*} r dz$$

and since:

$$\frac{\partial \widehat{\Pi}_1^3}{\partial x_1}(x_1, x_2^*) < 0 \text{ for } x_1 \in [x_2^* - 2m, x_2^* - m] \text{ and } m \leq \min\{a/[2(1-a)], 1/10\}$$

it is sufficient to observe that:

$$\lim_{m \rightarrow 0} \left[\Pi^* - \widehat{\Pi}_1^3(x_2^* - 2m, x_2^*) \right] = \frac{2+a}{16} > 0$$

2. {12} $x_1 \in [x_2^* - m, x_2^*]$:

$$\widehat{\Pi}_1^4(x_1, x_2^*) = \int_0^{x_1-m} p_{1.1}^m dz + \int_{x_1-m}^{x_2^*} r dz$$

and since:

$$\frac{\partial \widehat{\Pi}_1^4}{\partial x_1}(x_1, x_2^*) < 0 \text{ for } x_1 \in [x_2^* - m, x_2^*], r > 0 \text{ and } m \leq 1/10$$

it is sufficient to observe that:

$$\lim_{m \rightarrow 0} \left[\Pi^* - \widehat{\Pi}_1^4(x_2^* - m, x_2^*) \right] = \frac{2+a}{16} > 0$$

Finally, note that in the present case of $a \in [0, 1)$ the second-stage profit functions are continuous and hence, by Theorem 3 in Dasgupta and Maskin [1986a] a mixed-strategy price equilibrium always exists. Clearly, the corresponding profits may not exceed r . It follows that for m sufficiently small player 1 (and, by symmetry, player 2) cannot profitably deviate from the equilibrium characterized in Proposition 1. Finally, note that the candidate equilibrium locations are unique and hence any SPNE which entails playing the PSE where available must have locations (2.2). ■

Proof of Proposition 2. Again, some algebraic derivations are relegated to the on-line appendix, where numbers in curly brackets $\{\#\}$ represent a position therein where a particular statement is verified.

We begin by obtaining the “candidate” equilibrium locations (2.2). Clearly, for $z < x_1 - m$ the unique price equilibrium is $\{p_1 = x_2 - x_1, p_2 = 0\}$, i.e. firm 1 doing just enough to capture the whole market with firm 2 setting its price to zero. For $x_1 - m < z < z_1^m = \frac{1}{2}(x_1 + x_2 - 4m)$ firm 1 still attracts all consumers, earning a profit of $x_1 + x_2 - 2(z + m)$ {13}. Conversely, firm 2 takes the whole market for $z > z_2^m = \frac{1}{2}(2m + x_1 + x_2)$ {14}. For the intermediate $z_1^m < z < z_2^m$, either no equilibrium exists or there is a usual “competitive” equilibrium, obtained by finding the location of the “indifferent” consumer \tilde{x} and solving

$$\frac{\partial \{p_1(\tilde{x} - z)/m\}}{\partial p_1} = \frac{\partial \{p_2[1 - (\tilde{x} - z)/m]\}}{\partial p_2} = 0$$

for p_1, p_2 , which gives the equilibrium profit of firm 1:

$$\pi_c^* = \frac{[x_1 + x_2 - 2(z - m)]^2}{18m}$$

The necessary and sufficient condition for the price equilibrium to exist for all $z_1^m < z < z_2^m$ is $x_2 - x_1 \geq 2m$ {15}, which also ensures that $x_1 - m < z_1^m$. Assume first that $x_1 < m$. The expected profit of firm 1 is then:

$$\frac{1}{1 - m} \left[\int_0^{z_1^m} x_1 + x_2 - 2(z + m) dz + \int_{z_1^m}^{z_2^m} \pi_c^* dz \right]$$

which is increasing in x_1 for $x_1 < m$ and $x_2 - x_1 \geq 2m$ {16}. On the other hand, for

$x_1 \in [m, x_2 - 2m]$, the expected profit of firm 1 is given by:

$$\frac{1}{1-m} \left[\int_0^{x_1-m} x_2 - x_1 dz + \int_{x_1-m}^{z_1^m} x_1 + x_2 - 2(z+m) dz + \int_{z_1^m}^{z_2^m} \pi_c^* dz \right]$$

which evaluates to {17}:

$$\Pi_1(x_1, x_2) = \frac{4m^2 + 4mx_1 - 3x_1^2 - 4mx_1 + 2x_1x_2 + x_2^2}{4 - 4m}$$

By symmetry of the game, we have $\Pi_2(x_1, x_2) = \Pi_1(1 - x_2, 1 - x_1)$.

Since $\partial^2 \Pi_i / \partial x_i^2 = 3 / [2(m-1)] < 0$ {18}, the candidate equilibrium is given by {19}:

$$\frac{\partial \Pi_1}{\partial x_1} = \frac{\partial \Pi_2}{\partial x_2} = 0 \Leftrightarrow x_1^* = 1 - x_2^* = \frac{1}{4}(1 + 2m)$$

which means both players earn:

$$\Pi^* = \frac{3 - 12m + 28m^2}{16(1-m)}$$

and the condition $x_2^* - 2m \geq x_1^* > m$ is satisfied for $m \leq 1/6$ {20}.

We now establish the conditions for $\{x_1^*, x_2^*\}$ to form a SPNE, by showing that no deviation to a $x_1 > x_2^* - 2m$ can give player 1 an expected profit larger than Π^* for uncertainty big enough relative to a non-binding reservation price r .

First of all, the smallest such price is $r = x_2^*$. To see this, observe first that for any response to x_2^* , no consumer ever has to incur a total cost of more than x_2^* in any “competitive” price equilibrium {21}. Consequently, the associated equilibrium payoffs remain the same, while the profits from any price-deviation could only decrease with the reservation price imposed. This leaves any “competitive” price equilibria unaffected by the reservation price. On the other hand, in any “monopolistic” equilibrium the total cost incurred by a consumer cannot exceed the cost of travelling to the location of the player who sets his mill-price to zero. This means it cannot exceed r when the player located at x_2^* ends up with no market share. And when it is the deviating player who ends up with nothing, r does not matter for his second-stage equilibrium profit anyway.

The next step is to verify that with $r \geq x_2^*$, $x_1 \in (x_2^* - 2m, x_2^*)$, $p_2 = r$ and either firm located in $[z, z+m]$, firm 1 would choose a p_1 to capture the entire market subject to the imposed reservation price {22}. This gives an upper bound of $\bar{\pi}_l = (r - x_1 + z)$ on the second stage profits of the deviating firm when $x_1 - m < z < x_1 - m/2$ and

$\bar{\pi}_r = (r - m + x_1 - z)$ when $x_1 - m/2 < z < x_2$ {23}.

Since $x_2^* - x_1 < 2m$, for $x_1 - m < z < x_2^*$ we either have a “competitive” price equilibrium or none. Let:

$$\begin{aligned} z_1^u &= \frac{1}{8} \left[3 + 18m + 4x_1 - 6\sqrt{2m(12m + 4x_1 - 3)} \right] \\ z_2^u &= \frac{1}{8} \left[3 - 30m + 4x_1 + 6\sqrt{2m(12m + 4x_1 - 3)} \right] \end{aligned}$$

and it follows that firm 1 never wants to undercut the candidate equilibrium price of the rival for $z < z_1^u$, while the opposite is true for $z > z_2^u$ {24}.

1. Consider first the case of $x_1 \in (x_2^* - 2m, x_2^* - m)$. We then have {25}:

$$x_2^* > z_1^u > x_2^* - m > x_1 > z_2^u > x_1 - m$$

which means the price equilibrium does not exist for $z \in (x_1 - m, z_2^u)$ and for $z \in (z_1^u, x_2^*)$. Since $z_1^u > x_1 - m/2$ {26}, the best the deviating player can get in the latter case is $\bar{\pi}_r$. As for $z \in (x_1 - m, z_2^u)$, let $\hat{x}_1 = 3/4 + (5/2 - 3\sqrt{2})m$ and observe that $x_1 < \hat{x}_1 \Leftrightarrow z_2^u < x_1 - m/2$ {27}. Consequently:

- (a) for $x_1 \in (x_2^* - 2m, \hat{x}_1)$, the best the deviating player can get when $z \in (x_1 - m, z_2^u)$ is $\bar{\pi}_l$ and the maximum expected profit becomes:

$$\hat{\Pi}_1^a = \frac{1}{1-m} \left[\int_0^{x_1-m} (x_2^* - x_1) dz + \int_{x_1-m}^{z_2^u} \bar{\pi}_l dz + \int_{z_2^u}^{z_1^u} \pi_c^* dz + \int_{z_1^u}^{x_2^*} \bar{\pi}_r dz \right]$$

- (b) for $x_1 \in (\hat{x}_1, x_2^* - m)$, the best the deviating player can get is $\bar{\pi}_l$ when $z \in (x_1 - m, x_1 - m/2)$ and $\bar{\pi}_r$ when $z \in (x_1 - m/2, z_2^u)$, so that the maximum expected profit becomes:

$$\hat{\Pi}_1^b = \frac{1}{1-m} \left[\int_0^{x_1-m} x_2^* - x_1 dz + \int_{x_1-m}^{x_1-\frac{m}{2}} \bar{\pi}_l dz + \int_{x_1-\frac{m}{2}}^{z_2^u} \bar{\pi}_r dz + \int_{z_2^u}^{z_1^u} \pi_c^* dz + \int_{z_1^u}^{x_2^*} \bar{\pi}_r dz \right]$$

2. Suppose now $x_1 \in (x_2^* - m, x_2^*)$ and let $\tilde{x}_1 = 3/4 - m$. We have $z_1^u < z_2^u \Leftrightarrow x_1 > \tilde{x}_1$ {28} and:

- (a) for $x_1 \in (x_2^* - m, \tilde{x}_1)$ we have $x_1 - m/2 < z_2^u < z_1^u < x_2^*$ {29}, so that the maximum expected profit is still given by $\hat{\Pi}_1^b$.

(b) for $x_1 \in (\tilde{x}_1, x_2^*)$ we have $z_2^u > z_1^u$, so that no “competitive” equilibria exist and the maximum expected profit becomes:

$$\widehat{\Pi}_2^b = \frac{1}{1-m} \left[\int_0^{x_1-m} x_2^* - x_1 dz + \int_{x_1-m}^{x_1-\frac{m}{2}} \bar{\pi}_l dz + \int_{x_1-\frac{m}{2}}^{x_2^*} \bar{\pi}_r dz \right]$$

All in all, in order to ensure that no deviation from x_1^* is profitable, we need:

$$\forall x_1 \in (x_2^* - 2m, \hat{x}_1) : \Pi^* \geq \widehat{\Pi}_1^a \quad (\text{C1})$$

$$\forall x_1 \in (\hat{x}_1, \tilde{x}_1) : \Pi^* \geq \widehat{\Pi}_1^b \quad (\text{C2})$$

$$\forall x_1 \in (\tilde{x}_1, x_2^*) : \Pi^* \geq \widehat{\Pi}_2^b \quad (\text{C3})$$

and it can be shown {30} that the conditions (C1) - (C3) are all true for $m < \bar{m} \approx 1/9$ and $r \in [x_2^*, \phi(m)]$, where $\lim_{m \rightarrow 0} \phi(m) = +\infty$. By symmetry, the same conditions guarantee that no deviation from x_2^* is profitable for player 2.

It remains to recall that by Theorem 3 in Dasgupta and Maskin [1986b], a mixed-strategy price equilibrium exists in a Hotelling game with linear costs for any $x_1 < x_2$. Clearly, the expected profit of a firm in any such equilibrium cannot exceed the optimal monopoly profit ($\bar{\pi}_l$ or $\bar{\pi}_r$). Suppose the players play the pure-strategy equilibrium prices where available and mixed-strategy equilibrium prices otherwise. Then the first-stage expected profit of a player deviating from x_i^* cannot exceed the one based on getting pure-strategy equilibrium profits where available and the optimal monopoly profits otherwise. Hence, the subgame-perfect equilibrium as described in the Proposition exists and has locations (2.2). ■

Product differentiation decisions under ambiguous consumer
demand and pessimistic expectations

3.1 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 instance, Balvers and Szerb [1996] consider the effect of random shocks to the products' quality / desirability when the 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. Other authors, such as Aghion et al. [1993], focus on price-experimentation as an instrument for acquiring information about uncertain demand.

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 at the same time, 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 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, one can conjure a similarly intuitive reasoning to the opposite effect if the firms are pessimistic, in the sense of always trying to prepare for the worst-case scenario. In such case, an increase in the range of possible demand variations means a player has to consider a potentially larger strategic advantage on the competitor's behalf, especially when the players are highly specialized, leaving more room for shifts in consumer preferences to favour one over the other. To insure against this threat, a pessimistic player might then want to make her product more similar to that of the rival when uncertainty increases.

A particularly useful illustration of the described mechanism is offered by the online sports betting industry. In general, the betting market is associated with a large degree of product differentiation, as different bookmakers specialize in different sports, types of bets and outcomes of a particular sporting event (for instance, some are known to offer better prices on the favorites, others on the underdogs of a competition). This gives them flexibility in balancing the bets, while avoiding head-to-head competition in terms of the overall house edge¹, which a 'recreational' bettor will find difficult to compare between bookmakers with different relative prices, particularly for events with several possible outcomes, such as horse races. Those events are characterized by a significant increase of the uncertainty about bettors' preferences shortly before the start of the competition, as evidenced by intensified trade and increased price volatility in betting exchanges (see Smith, Paton, and Williams [2009] for more information). Interestingly, this coincides with the odds baskets offered by different bookmakers ceasing to be differentiated (with a smaller house edge), i.e. everyone quotes exactly the same price on every horse.

The explanation might lie in the fact that bookmakers are unlikely to know the exact distribution of the punters' betting preferences, although they might be able to place them within a certain range (for instance, an outsider will not suddenly become the clear

¹i.e. the percentage of bettors' money the bookmaker is aiming to keep as revenue regardless of the outcome of the event

favorite of a race). Furthermore, a unique feature of the gambling industry is its extremely pessimistic approach to uncertainty, manifested by the traditional objective of ‘balancing the books’, i.e. effectively focusing on the worst-case scenario in which the outcome that attracted the largest volume of bets is realized. Assuming the bookmakers extend this attitude to uncertainty about consumer preferences, an increase in the range of possible demand variations will make it more undesirable for any particular firm to differentiate its odds from those of the competitors, because there is more room for the betting patterns to shift ‘against it’, in the sense of the books becoming more unbalanced due to bettors switching to and from other bookmakers in search of better odds. Consequently, losses associated with the success of the most excessively backed contestant are becoming more severe.

The existing literature outlined earlier is unable to accommodate the above mechanism, because of its reliance on the common prior assumption in modelling demand uncertainty. On the one hand, it seems reasonable to assume that the firms are not completely certain of the exact consumer preferences at the time of designing the product or choosing the location of their outlet. On the other hand, for all firms to unanimously form precise probabilistic estimates of all potential demand realizations may be too much to ask of an industry. Thus, the present paper considers an altogether different scenario, where the firms are ignorant of the distribution of demand fluctuations and possibly differ in their resolution of the resulting ambiguity.

The setting is the same as in MZ, except the demand is allowed to vary not only in location, but also in its price elasticity, as captured by the transportation cost parameter². In addition, the firms know only the support of the distribution of those changes, which necessitates a different payoff specification for the reduced location game. Instead of calculating the expected value of the second stage Nash Equilibrium profits for a particular location-pair, the firms consider a weighted average of the highest and lowest of those profits, i.e. use the Arrow and Hurwicz [1972] α -maxmin criterion to resolve the ambiguity³. The definition of uncertainty must also change. In MZ it was specified as the variance of the distribution of the shock shifting consumer preferences, which is not applicable in the absence of a common prior. Instead, an increase of demand uncertainty will be modelled *via* spreading the support of the random demand fluctuations, which is similar to the approach taken in Meagher and Zauner [2005]. The terms ‘ambiguity’ and ‘uncertainty’

²The 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, which is what is henceforth meant by price elasticity.

³For the model discussed here, the α -maxmin profits coincide with the α -Maxmin Expected Utility (as in Ghirardato, Maccheroni, and Marinacci [2004]) of a risk-neutral agent.

will be used interchangeably in the context of the present model, while ‘uncertainty’ will be exclusive to MZ.

It turns out that an increase of uncertainty about the demand’s location decreases the equilibrium product differentiation (and with it, the resulting second-stage prices and profits) when the firms are sufficiently pessimistic, in the sense of assigning a high enough weight to the worst-profit scenario. This is because a pessimistic player effectively assumes that the consumer preferences will move in a way offering a strategic advantage to the counterpart ahead of the second-stage price competition. By locating closer to the rival, he partly insures against this possibility, because even if the customers find the product of the other firm more suitable, his own one, being similar, is not so badly handicapped.

Surprisingly, the effect of pessimistic expectations is moderated by uncertainty about the price elasticity of demand (represented by a possibility of lower transport costs) despite the fact that the transportation cost parameter has no effect on location decisions in the Hotelling framework. This is because a pessimist will expect to see unfavorable consumer preferences combined with competitive pricing due to high elasticity of demand. A possibility of lower costs will make the price competition in such worst-case scenario even more intense, so that the pessimistic outcome of the uncertainty becomes even more threatening. However, it also becomes more costly to insure against this threat, because a larger reduction of product differentiation is required in order to achieve the same second-stage profit improvement in those unfavorable circumstances, resulting in a more significant reduction of one’s strategic advantage in case of the demand being favorably located. Consequently, a firm must now be more pessimistic in order to continue to decrease product differentiation in response to an increase of uncertainty about the demand’s location.

The results are, to a large extent, robust to a change in timing, such that the pricing decisions are made before the resolution of the uncertainty. Despite the existence of multiple price equilibria for ‘not too asymmetric’ locations, both the highest and lowest possible equilibrium prices are decreasing in demand location uncertainty when the firms are pessimistic, despite the effect being dampened by uncertainty about transport costs. Thus, uncertainty about the placement of consumer demand still makes pessimistic producers more competitive, although less so when faced with a possibility of a highly price elastic demand.

Finally, whenever a particular firm adapts a more pessimistic approach, its equilibrium location is further towards the competitor’s end of the market, with the rival withdrawn into his own hinterland. As a result, the pessimistic firm becomes better off in equilibrium at the counterpart’s expense, regardless of the eventual demand realisation, i.e. of whether or not the firm’s pessimistic expectations prove justified. This suggests that if the

duopolists could commit themselves to an attitude of ‘preparing for the worst’, then they might use pessimism as a form of strategic deterrence, preventing the competitor from targeting their own market niche. Alternatively, rather than through deliberate commitment, a pessimistic approach towards uncertainty could become prevalent through the elimination of firms who fail to adapt it and do worse relative to the competitors.

However, this creates a prisoner’s Dilemma situation, because when *both* firms become more pessimistic, product differentiation decreases and *both* fall victim of the intensified price competition in the second stage of the game. In a sense, the players’ self-imposed pessimism becomes a self-fulfilling prophecy.

As an alternative, Section 3.4 discusses a possibility of the same competitive outcome being achieved by means of corporate tax rebates. Thus, it may be concluded that the presence of ambiguously distributed demand variations leads to the product market being more competitive.

3.2 The Model

As indicated above, the setting is, in general, the same as in 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 stage. 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 $[-L, L] \times [t_0, 1]$, where $L \in [0, \frac{1}{2}]$ and $t_0 \in (0, 1]$ ⁴. The difference from MZ is introducing uncertainty about transportation costs (MZ assumes $t = 1$), as well as the fact that the exact probability distribution of (M, t) is unknown and so is the expected value of the second stage Nash Equilibrium profits. Instead, the players’ payoffs in the reduced location game are given by a weighted average of the lowest and highest possible profits, i.e. are computed using the Arrow/Hurwicz α -maxmin criterion instead of the expected value. More specifically, let $\pi_i^*(x_1, x_2, M, t)$ be the second-stage unique Nash Equilibrium profit associated with a particular location-pair and demand realization.

⁴The assumption that transportation costs are always no greater than 1 can be imposed without loss of generality. The assumption $L \leq 1/2$ was imposed in MZ for the purpose of mathematical tractability and is equally useful here.

Then the first-stage payoffs are given by:

$$\begin{aligned} \Pi_i(x_1, x_2) = & \alpha \left[\min_{(M,t) \in [-L,L] \times [t_0,1]} \pi_i^*(x_1, x_2, M, t) \right] + \\ & + (1 - \alpha) \left[\max_{(M,t) \in [-L,L] \times [t_0,1]} \pi_i^*(x_1, x_2, M, t) \right] \end{aligned}$$

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

3.3 Results

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

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

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, we have:

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

Note that a player can always ensure a positive α -maxmin profit, by locating at $x_i = -x_{-i}$, so that $\bar{x} = 0$ and $(M - \bar{x}) \in (-3/2, 3/2)$ for both $M = -L$ and $M = L$ (recall $L < 1/2$ by assumption), i.e. there is a competitive equilibrium in both the highest-profit and lowest-profit scenarios.

We will now show that any Nash Equilibrium of the reduced location game must satisfy $(M - \bar{x}) \in (-3/2, 3/2)$ for $M \in \{-L, L\}$, i.e. that it must result in the best and worst second-stage equilibria being competitive. To this end, consider the following cases:

1. Player 1 captures the entire market for $M = -L$, while player 2 does the same for $M = L$. 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 both $M = -L$ and $M = L$. This cannot constitute a Nash Equilibrium of the reduced location game, because the other player is able to improve on her zero payoff.
3. Player 1 monopolizes the market in the highest-profit scenario of $M = -L$ and a competitive equilibrium follows for $M = L$. The α -maxmin profit of player 2 then equals:

$$\Pi_2^{0c}(x_1, x_2) = \alpha \times 0 + (1 - \alpha)(x_2 - x_1)[3 + 2(L - \bar{x})]^2 / 18$$

For $\alpha = 1$ the above is equal to 0, so that player 2 benefits from re-locating to $x_2 = -x_1$ and ensuring competitive equilibria for $M \in \{-L, L\}$. As for the case of $\alpha < 1$, differentiating the above expression with respect to x_2 gives:

$$(1 - \alpha)(3 + 2L + x_1 - 3x_2)[3 + 2(L - \bar{x})]$$

we have $3 + 2(L - \bar{x}) > 0$, since $(L - \bar{x}) \in (-3/2, 3/2)$. The only stationary point left is therefore $x_2^s = 1 + (2L + x_1) / 3$, so that $\bar{x} = (3 + 2L + 4x_1) / 6$ and the monopolistic equilibrium at $M = -L$ implies:

$$L + (3 + 2L + 4x_1) / 6 \geq 3/2 \Leftrightarrow x_1 \geq \underline{x}_1 = (3 - 4L) / 2$$

The α -maxmin profit of player 1 equals:

$$\Pi_1^{cm}(x_1, x_2) = \alpha t_0 (x_2 - x_1) [2(L - \bar{x}) - 3]^2 / 18 + (1 - \alpha)(x_2 - x_1) [1 + 2(L + \bar{x})]$$

differentiating with respect to x_1 and substituting x_2^s for x_2 we obtain:

$$\begin{aligned} \phi(x_1) = & -\frac{16\alpha t_0}{81} x_1^2 + \left(\left[2 + \frac{4(8L - 15)t_0}{81} \right] \alpha - 2 \right) x_1 + \\ & + \frac{2[81L - 2(L - 3)(4L - 3)t_0] \alpha}{81} - 1 - 2L + \alpha \end{aligned}$$

which is quadratic concave with a maximum at:

$$x_1 = ([81 + 2(8L - 15)t_0] \alpha - 81) / 16t_0 \alpha < \underline{x}_1$$

Furthermore:

$$\phi(x_1) = 2L - 4(1 - \alpha) - (2L + [2 + L(16L/9 - 4)] t_0) \alpha$$

which is negative under the assumptions on α , L and t_0 . Hence, the derivative $\phi(\cdot)$ is negative, i.e. whenever player 2 is satisfied with her current location, player 1 wants to move leftward. Consequently, there can be no location equilibrium with player 1 monopolizing the market for $M = -L$ and a competitive equilibrium at $M = L$. Similarly for the opposite case of player 2 monopolizing the market for $M = L$ and a competitive equilibrium at $M = -L$.

The only remaining possibility is that the price equilibrium is competitive for $M \in \{-L, L\}$, in which case payoffs equal:

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

The corresponding first order conditions are:

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

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

This has three possible solutions, two of which fail to satisfy the competitive equilibrium condition $(M - \bar{x}) \in (-3/2, 3/2)$ for $M \in \{-L, L\}$ (see online appendix). The remaining solution is:

$$x_1^* = -x_2^* = \frac{(3 + 2L)^2 (1 - \alpha) + (3 - 2L)^2 t_0 \alpha}{4 [2L((t_0 + 1)\alpha - 1) - 3 - 3(t_0 - 1)\alpha]} \quad (3.1)$$

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 $x_1 < x_2^*$ and monopolize the market for $M = L$, since this requires:

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

while we have:

$$3 + 2L - x_2^* < x_2^* \Leftrightarrow \frac{(9[1 - t_0] + 4L[3 + L + 3(L - 1)t_0])\alpha - (3 + 2L)^2}{2L[(t_0 + 1)\alpha - 1] - 3[1 + (t_0 - 1)\alpha]} < 0$$

where both the numerator and the denominator of the above fraction are negative under the imposed parameter assumptions.

Player 1 can still deviate to a $x_1 < x_2^*$ such that he will monopolize the market for $M = -L$ only. However, it turns out that the first derivative of Π_1^{cm} with respect to x_1 is negative for $x_2 = x_2^*$ and $L + (x_1 + x_2^*)/2 > 3/2$. Similarly, when it is player 2 who monopolizes the market for $M = L$ only, i.e. when $L - (x_1 + x_2^*)/2 > 3/2$ and $-L - (x_1 + x_2^*)/2 \in (-3/2, 3/2)$, then the derivative of:

$$\Pi_1^{0c}(x_1, x_2^*) = \alpha \times 0 + (1 - \alpha)(x_2^* - x_1)[-3 + 2(-L - (x_1 + x_2^*)/2)]^2 / 18$$

with respect to x_1 is positive. 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 $\partial \Pi_1^{0c}(x_1, x_2^*) / \partial x_1$ 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 the online appendix for details of all the algebraic derivations).

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).

We now turn to the comparative statics of the above solution. Firstly, for symmetric locations the equilibrium product differentiation is $\Delta^* = 2x_2^*$, which is decreasing in α , t_0 and:

$$\begin{aligned} \frac{\partial \Delta^*}{\partial L} &= \frac{[(t_0 - 1)\alpha + 1] [(3 + 2L)^2 (1 - \alpha) - (3 - 2L)^2 t_0 \alpha]}{(3 + 3(t_0 - 1)\alpha - 2L [(t_0 + 1)\alpha - 1])^2} < 0 \Leftrightarrow \\ &\Leftrightarrow \alpha > \hat{\alpha} = \frac{(3 + 2L)^2}{(3 + 2L)^2 + (3 - 2L)^2 t_0} \end{aligned}$$

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^* .

The above results may be summarized as follows:

Proposition 5 *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 -maxmin value of the ex-post second-stage equilibrium profits. Then the unique equilibrium locations are given by (3.1), while the corresponding product differentiation, as well as the second-stage equilibrium prices and profits, are all decreasing in L for α sufficiently large.*

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 realisation, 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, Teo, and Wang [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.

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:

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

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.

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 $L \in [0, 1/2]$ ($\alpha > 4/[4 + t_0]$).

Returning to the sports betting example invoked in the introduction, it would appear that the bookmakers' degree of pessimism is large even relative to a considerably wide spectrum of possible demand variations. For this reason, an increase of demand uncertainty shortly before the start of a race makes them reduce the differentiation in the offered baskets of odds, despite the fact that this brings about a more competitive house edge. This is because any bookmaker who chose to offer different odds from the rest would face a particularly bad worst-case scenario, in which the betting patterns shift significantly in a way favorable to the competitors, making his bets extremely unbalanced. In order to insure against this threat, it is better to offer the same product as everybody else, because without customers 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 disadvantaged⁶.

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

⁶It may also be noted that no bets are accepted by traditional bookmakers after the start of the race, despite this form of betting being very popular in the betting exchanges. This may be due to the practical difficulties associated with coordinating their odds at a stage when betting patterns change within seconds.

a certain realisation 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. Thus, uncertainty about the intensity of price competition 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 the intensified 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 ($L < 1/2$), 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.

Finally, it is interesting to consider a measure of product differentiation relative to $2L + 1$. In the present framework, this turns out to be decreasing in the size of the uncertainty for any $\alpha \in [0, 1]$, i.e.:

$$\frac{\partial [(x_2^* - x_1^*) / (2L + 1)]}{\partial L} < 0$$

(see online appendix for the full calculation)

The above means that as uncertainty increases, the firms are located closer together relative to the spread of possible consumer tastes, which range from $(-L - 1/2)$ to $(L + 1/2)$. This occurs regardless of the players' degree of pessimism, which is only essential to ensure that Δ^* also decreases in absolute terms. Such would be the conclusion if uncertainty was modelled as in Krol [2011], based on a varying length of a consumer cluster, to be located within a *fixed* space of feasible tastes. All the remaining comparative statics results, including the ones in the following section, would then coincide with the current specification.

3.4 Extensions

We now turn to consider the possibility of the players being characterized by different degrees of pessimism. One would strongly expect the comparative statics results of the previous section to continue to hold, i.e. an increase of uncertainty should still decrease the equilibrium product differentiation for *both* duopolists not too optimistic relative to the minimum transportation costs. For this reason, let $L = 1/2$ and $t_0 = 1$, with the focus instead on the effect of a *ceteris paribus* change of attitude by a particular player on the equilibrium profits. The first-stage payoff function is:

$$\Pi_i(x_1, x_2) = \alpha_i \pi_i^*(x_1, x_2, -[-1]^i/2, 1) + (1 - \alpha_i) \pi_i^*(x_1, x_2, [-1]^i/2, 1)$$

The logic of the proof derives from that presented in the previous section. Firstly, the discussion of cases (1) – (3) in the previous section applies just as well to the present situation, i.e. the players' best-response mapping still ensures that any equilibrium locations must be such that the resulting price equilibrium is competitive for any realisation of the uncertainty. Consequently, the result is, once again, obtained as the unique solution to the first order conditions within the range of qualifying locations. However, because of the complexity of the involved algebra, only a sketch of the proof is provided here, with the details of all derivations (as well as the exact solution formulae) relegated to the online appendix.

Let $x_i^*(\alpha_i, \alpha_{-i})$ denote the unique equilibrium locations, where uniqueness 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.

I will now show that a *ceteris paribus* increase of α_i increases the ex-post equilibrium profit of player i and decreases that of the other player, for any demand realisation $M \in [-1/2, 1/2]$. To this end, consider the ex-post competitive equilibrium profit of player 1 given a particular value of M :

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

differentiating with respect to α_1 gives a product of:

$$[-3 + 2M - x_1^*(\alpha_1, \alpha_2) - x_2^*(\alpha_2, \alpha_1)] / 18 < 0$$

and:

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

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 that $\partial x_1^*(\alpha_2, \alpha_1) / \partial \alpha_1 > \partial x_2^*(\alpha_2, \alpha_1) / \partial \alpha_1 > 0$, 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 (see Figure 3.1). 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:

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

which can be shown to be the case for all $\alpha_1, \alpha_2 \in [0, 1]$. A converse of this argument holds for a change in the attitude of player 2 (by symmetry of the game) and we may summarize these findings as follows:

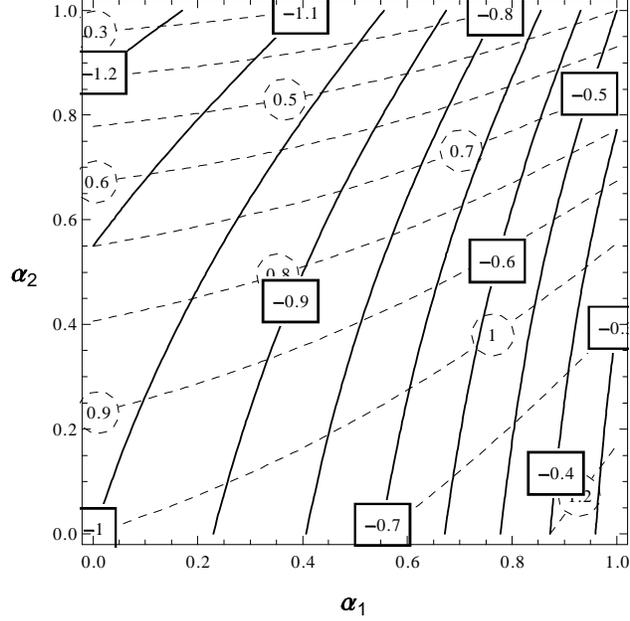


Figure 3.1: Iso-location curves of Player 1 and 2 (dashed), depending on the respective degrees of pessimism α_i . When one firm becomes more pessimistic, both shift towards the other firm's end of the market, but product differentiation decreases.

Proposition 6 Consider a variant of the Hotelling duopoly game in which the distribution of the median of the (uniform) consumer preferences M is unknown with support on $[-1/2, 1/2]$. Suppose the firms choose locations so as to maximize the α -maxmin value of the ex-post second-stage equilibrium profits, based on their respective degrees of pessimism $\alpha_1, \alpha_2 \in [0, 1]$. Then the equilibrium locations $x_i^*(\alpha_i, \alpha_{-i})$ are unique up to symmetry and such that for any $i, j \in \{1, 2\}$:

- (1) : $(-1)^i [\partial x_i^*(\alpha_i, \alpha_{-i}) / \partial \alpha_i] < (-1)^i [\partial x_{-i}^*(\alpha_{-i}, \alpha_i) / \partial \alpha_i] < 0$
- (2) : $\forall M \in [-1/2, 1/2] : (-1)^{i-j} \partial \pi_i(x_1^*(\alpha_1, \alpha_2), x_2^*(\alpha_2, \alpha_1), M) / \partial \alpha_j > 0$

This result is particularly interesting if the firms have some way of committing to a pessimistic policy, for instance, by appointing cautious CEO's or by putting themselves in a position where losing customers due to a sudden change in preferences could mean bankruptcy, thereby making it necessary to take the necessary precautions. In such case, 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 way, the firms' self-imposed pessimism becomes a self-fulfilling prophecy.

Alternatively, rather than through conscious commitment, 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. The corresponding product differentiation is equal to 1, less than the one resulting from the players following *any* common prior. Since the associated prices are also lower, the presence of ambiguity benefits the consumers, although the average transport costs they pay (and hence, the socially-optimal locations) will depend on the actual distribution of M . In contrast, the presence of ambiguously distributed demand fluctuations adversely affects the firms (compared with the certainty case), the opposite of what happens when the demand variations follow a commonly known pattern (as demonstrated by the MZ model).

Interestingly, the equilibrium locations associated with $\alpha_1 = \alpha_2 = 1$ and maximum uncertainty ($L = \frac{1}{2}$) are $x_2^* = -x_1^* = \frac{1}{2}$, and coincide with the ones which are socially-optimal, given customers uniformly distributed on the $[-1, 1]$ interval (see, for example Gabszewicz and Thisse [1992] or Lederer and Hurter [1986]). In the current framework, all that is known to the firms is that the consumer preferences are contained in $[-1, 1]$. Hence, the equilibrium locations seem like a sensible choice for an equally uninformed social-planner.

In fact, a planner could achieve the same competitive outcome by means of taxation, even when the firms' attitudes are not subject to change. To see this, note first that taxing profits would not affect the second-stage optimal price response (to the counterpart's price), so long as the amount net of tax is increasing with the gross profit. Hence, the unique Nash Equilibrium prices remain unchanged. However, the same cannot be said about the equilibrium location decisions under uncertainty. In particular, consider a policy that entails a full tax rebate on amounts up to the worst-case second-stage equilibrium profit given the most 'pessimistic' equilibrium locations, i.e. up to $\pi_i(x_1^*(1, 1), x_2^*(1, 1), -(-1)^i L)$, and an arbitrarily high proportional tax thereafter. Locations $x_i^*(1, 1)$ maximize the worst-case profit, so a unilateral relocation must decrease it. Even if it improves the *highest* possible profit, any such improvement is subject to tax. Consequently, a sufficiently high tax rate (relative to the players' degrees of optimism) will ensure that the improvement fails to offset the decrease in the worst-case amount in terms of the α -maxmin profit. As a result, the highly competitive locations $x_i^*(1, 1)$ can be implemented as a Subgame-Perfect Nash Equilibrium regardless of the players' actual degrees of pessimism. Indeed, a reasoning similar to the one above may also be applied to imposing a tax policy in the presence of a common prior.

Finally, 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 realisation 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. 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, we have:

$$\pi_i(x_1, x_2, p_1, p_2) = \alpha \left[-L + 1/2 - (-1)^i \tilde{x} \right] p_i + (1 - \alpha) \left[L + 1/2 - (-1)^i \hat{x} \right] p_i$$

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 are not 'too asymmetric', so that they satisfy:

$$x_1 + x_2 \in [-x_T, x_T], \quad x_T = (1 - 2\alpha) [(4\alpha - 2)L - 1] (1 - t_0) / (1 + t_0) \quad (3.2)$$

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.

Because of the possible multiplicity of price equilibria, the firms' location decisions will not be examined here. Nevertheless, assuming that locations are indeed not 'too asymmetric' (in the sense explained above), it may be observed that:

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

leading to the following statement:

Proposition 7 *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' locations satisfy condition (3.2) and that the firms choose prices so as to maximize the α -maxmin value of the resulting profits. In any Nash Equilibrium, the firms set identical prices, where the lowest and highest possible equilibrium prices 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, the results are in line with those in the previous section. 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 can be dampened by a possibility of a more price elastic consumer demand.

Furthermore, we observe a tendency of the firms to mimic the competitor, where in the present case this takes an extreme form of what is effectively a coordination problem, even when each firm's initial situation ahead of the pricing stage is different. Once again, this may be related to the sports betting example, where the odds on offer converge shortly before the start of a race. Interestingly, some of the major bookmakers offer a so called 'best odds guarantee'⁷ within a given time before the start of the biggest races, pledging to at least match the competitors' prices when paying out the winnings. This could be seen as an attempt to solve the coordination problem, because the effective prices are all equal to the maximum of the ones actually offered by the bookmakers participating in the scheme.

3.5 Concluding Remarks

The paper examined a variation of the "Product differentiation and location decisions under demand uncertainty" model by Meagher and Zauner, in which the firms are unaware of the exact distribution of demand fluctuations, but resolve the resulting ambiguity using the Arrow-Hurwicz α -maxmin criterion. The change made it possible to accommodate a scenario in which firms approximate the competitors' behaviour in response to an increase of demand uncertainty, as illustrated by the sports betting industry.

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 respond to an increase of uncertainty by making his offer more similar to that of the other firm, thereby leaving less room for his product being disadvantaged. This mechanism was analyzed in the context of the current model, where it turns out that, contrary to the existing literature, an increase of uncertainty about the location of consumer preferences makes the equilibrium more competitive when the firms are sufficiently pessimistic. In particular, product differentiation decreases, and with it the second-stage equilibrium prices and profits for any realisation of the uncertainty. This effect is moderated by

⁷see, for example, <http://tinyurl.com/2g65uyy>

uncertainty about the transport cost parameter (reflecting the price elasticity of demand), which has no effect on location decisions under certainty, but interacts with uncertainty about the placement of consumer demand.

When price-competition takes place before the realisation of uncertainty, it is affected in a similar way. In particular, despite the existence of multiple equilibria for ‘not too asymmetric’ locations, the highest and lowest possible equilibrium prices both decrease when demand location uncertainty increases, as long as the firms are pessimistic. Once again, the positive effect of this type of uncertainty on the degree of competition between pessimistic firms is moderated by a possibility of higher price elasticity of demand.

Finally, whenever a particular firm adapts a more pessimistic approach, it becomes better off at the competitor’s expense, suggesting that ‘being prepared for the worst’ could serve as a form of strategic deterrence. Whether by conscious changes in approach, or *via* elimination of underperforming firms, pessimistic attitudes towards ambiguity are likely to become prevalent, or else the corresponding outcome may still be implemented by means of taxation. This entails considerably less strategic product differentiation and lower prices than under certainty, suggesting that ambiguously distributed demand variations make the market more competitive, as opposed to ones that can be characterized by a common prior.

Appendix

Click on the link <http://tinyurl.com/36pnr14> to download (Wolfram Mathematica file) or enter directly into browser. A free Wolfram Mathematica Reader software is available for download at: <http://www.wolfram.com/products/player/>, and a non-interactive PDF version of the appendix is available at: <http://tinyurl.com/3nlwnzt>.

Part II

Homogeneous Goods

4.1 Introduction

Game theoretic modelling of economic phenomena typically entails simplifying the investigated strategic interactions by imposing restrictions on the agents' strategy sets and focusing only on selected aspects of the problem at hand. This is so as to ensure tractability of the model and possibly determine the precise outcome of the game *via* the Nash Equilibrium concept or its later refinements.

Nevertheless, the process of restricting the set of choices available to each player can be carried through in many alternative ways, subject to different, often equally plausible justifications. Naturally, in the absence of formal model selection criteria, the theorist's choice is bound to gravitate towards specifications leading to a unique equilibrium, ideally one with comparative statics properties consistent with the stylized facts. It is then quite surprising that so little attention appears to have been given in the literature to formulating such objective criteria. Compared with the abundance of rigorous non-cooperative game theory solution concepts, the setup of problems to which they are applied seems somewhat arbitrary. The present paper is an attempt to fill this gap.

To this end, I put forward a 'stable convention' criterion, based on the idea that the structure of a strategic interaction must be determined, and become accessible to the players, prior to the formation of conjectures about the counterparts' play, i.e. under a form of strategic uncertainty. Thus, a restriction of the players' strategy sets is said to constitute a stable convention if it is sustainable based only on the knowledge that one's counterparts

will adhere to the convention and choose *some* strategies within their restricted sets. More precisely, a convention is stable when it complements the set of strategies which are irrational, in the sense of being weakly dominated by some conventional strategy, conditional on others playing in a conventional manner.

In fact, every stable convention is also a von Neumann - Morgenstern (vN-M) stable set, for the underlying dominance relation specified accordingly. In this sense, the present notion constitutes an application (as a model evaluation criterion) of a recent generalization of the vN-M concept proposed by Luo [2001] and further axiomatized in Luo [2009], in which the dominance relation is conditional on the set of available alternatives. The postulated criterion also generalizes the strict Nash Equilibrium, in that the latter is a singleton stable convention. Finally, it is related, albeit less directly, to other set-valued concepts, such as strategic stability by Kohlberg and Mertens [1986], rationalizability by Bernheim [1984] and Pearce [1984], as well as the notion of sets closed under rational behaviour (Basu and Weibull [1991]), but should not be confused with the (stochastically) stable convention in the sense of Young [1993]. Those relationships are discussed in the following section of the paper, though an important caveat to be mentioned at this point is that the idea put forward here, unlike the above notions, is not intended as a solution concept, but rather as a supplementary evaluation / comparison criterion for alternative model specifications postulated for other reasons.

To illustrate and motivate the proposed notion, in Section 4.3 it is used to revisit the Cournot vs. Bertrand debate in oligopoly theory. On the one hand, Cournot equilibria ‘exist in a wide range of circumstances and exhibit empirically plausible properties’ (Vives [1989])¹, as opposed to the Bertrand framework, which typically yields either the well-known paradoxical result (see, for e.g. Harrington [1989]), or an indeterminate outcome, as in Dastidar [1995]². On the other hand, the Cournot model was often described as being ‘right for the wrong reasons’ (Fellner). Although it leads to an outcome consistent with the stylized facts, it does so without allowing for any pricing decisions, whereas in reality firms are at liberty to set both quantities and prices.

Attempts to advocate the Cournot framework in the face of such criticism usually stem from appreciating that exogenous institutional or technological factors, e.g. the general cost structure or relative flexibility of prices vs. quantities, may determine the manner of competition between firms. The aim is then to demonstrate that the outcome of Cournot competition is consistent with equilibrium play in games with richer strategy spaces under some circumstances. Unfortunately, a pure strategy Nash Equilibrium may not exist in

¹I will not attempt to review the vast Cournot literature here, but see, for instance, Novshek [1985] for existence results, and Dixit [1986] or Novshek [1980] for comparative statics.

²See Saporiti and Coloma [2010] for a compact but exhaustive summary of the Bertrand literature.

price-quantity games when the production costs are ‘sunk’ (Friedman [1988]), i.e. incurred upon offering the goods for sale, rather than ‘sales-dependent’, i.e. based on the quantity of output that is eventually sold, as in Dixon [1992]. A similar indeterminacy of results, but due to *multiplicity* of equilibria, is prevalent in ‘supply-function’ models, where firms simultaneously commit to price contingent supply schedules (see Grossman [1981]). However, as demonstrated by Klemperer and Meyer [1986, 1989], the indeterminacy problem is alleviated by introducing uncertain consumer demand. In particular, the equilibria that survive in the presence of uncertainty exhibit Cournot-type (quantity) strategies when the marginal costs of production are relatively steep.

Another line of research, initiated by Kreps and Scheinkman [1983], imposes more structure on the game with respect to the timing of moves: firms build production capacities in the first stage of the game, before competing in prices subject to the imposed capacity constraints. It turns out that the Cournot outcome constitutes a subgame-perfect equilibrium of the game, in the sense that firms choose the equilibrium Cournot quantities as first stage capacities, and then charge the demand price associated with their aggregate production capacity. Although Davidson and Deneckere [1986] demonstrated the dependency of the last result on a particular (‘efficient’) rationing rule, Madden [1998] generalized the result in this respect by assuming production costs are ‘sunk’, while Acemoglu et al. [2009], as well as Moreno and Ubeda [2006], presented modifications of the model which make rationing altogether irrelevant, by introducing inelastic demand and reservation prices respectively.

The last innovation improves on the Kreps and Scheinkman result, in that a pure strategy second-stage price equilibrium is no longer absent for some capacity choices. What makes this framework particularly relevant here, is that by offering up to the selected capacity of output at at least the specified reservation price, firms choose their strategies out of a class of supply functions. Crucially, this class is sufficiently general to accommodate the complete set of either Cournot or Bertrand outcomes, depending on how the strategy sets are further restricted. Thus, while being relatively simple, the setting is well-suited for comparing the two alternative competition modes by means of the proposed model evaluation criterion.

Applying it to a single-stage version of the Moreno and Ubeda game makes it possible to argue in favour of Cournot without reference to dynamics. In particular, it turns out that a restriction of the initial strategy sets equivalent to the classic Cournot game is a stable convention if and only if production costs are high relative to the number of firms and difficult to recover for unsold output. In contrast, Bertrand competition is never stable. Thus, the entire Cournot game, rather than just its equilibrium outcome, may be

justified as a valid simplification of oligopolistic competition when players set reservation, rather than exact prices, an assumption applicable to many modern-day trading platforms (see Section 4.3.1). Expecting others to behave in *some* way compliant with Cournot competition is sufficient to make it irrational to unilaterally violate this convention. In this sense, when the derived conditions are met, the Cournot model is right for the very right reasons.

4.2 The concept

4.2.1 Definition and basic properties

Consider a normal-form game $G = (N, \{S_i\}_{i \in N}, \{U_i\}_{i \in N})$, where $N = \{1, 2, \dots, n\}$ is the set of players, S_i is the set of (mixed) strategies of player i and $U_i(s) \in \mathbb{R}$ is the associated payoff when a strategy profile $s \in \times_{j \in N} S_j$ is played. Also let $S_{-i} \equiv \times_{j \in N \setminus \{i\}} S_j$ be the set of strategy profiles of players other than i .

Definition 8 *A collection of strategy sets $\{S'_i\}_{i \in N}$ is a **stable convention** when for each $i \in N$ we have $S'_i \subseteq S_i$ and when it is true that $s_i \in S_i \setminus S'_i$ if and only if there exists a $s'_i \in S'_i$ such that $U_i(s'_i, s'_{-i}) \geq U_i(s_i, s'_{-i})$ for all $s'_{-i} \in S'_{-i}$, where the inequality is strict for some s'_{-i} .*

In other words, no **conventional** strategy of any player i (one in S'_i) is weakly dominated by another conventional strategy against the set S'_{-i} of other players' conventional strategy profiles, while every unconventional strategy (i.e. not in S'_i) is weakly dominated by some conventional strategy.

This is based on the premise that players determine what sort of game they are due to play *a priori*, or independently of how they and their counterparts would choose to play it. Since players make conjectures about the rivals' choices based on their knowledge of the game, including the sets of strategies that could be played, the structure of the game should be decided prior and without reference to any conjectures about its exact outcome. In this sense, a convention is stable when it is best for everyone to adhere to it as long as others are believed to do so, even when nothing more specific is known about their actions.

Consider then a player who gets to choose between two options, one of which is either equally good or strictly better than the other, but is completely ignorant of which of those possibilities is the case. Naturally, it is only rational to choose the first option. Similarly, strategies weakly dominated by conventional strategies against the counterparts' conventional play will never be used, and must be the ones that lie outside the convention if the latter is to be self-sustainable based on rational behaviour.

Example 9 Consider the following two-player normal-form game:

	D	E	F
A	(0, 6)	(2 , 5)	(1 , 6)
B	(4, 2)	(0, 4)	(3, 0)
C	(1, 1)	(1 , 5)	(4 , 3)

There exists no Nash Equilibrium in pure strategies, and no strategies can be eliminated as (weakly) dominated or non-rationalizable. However, if we eliminate B, then D is weakly dominated by F, and if we eliminate D, then B is strictly dominated by C. Neither of the remaining strategies is then weakly dominated by another, i.e. removing both B and D results in a stable convention.

It is worth noting that, in the above example, the stable convention $\{\{A,C\}, \{E,F\}\}$ is not ‘closed under rational behaviour’ (CURB) in the sense of Basu and Weibull [1991], because D is a (weak) best response to A. This would not be the case if a CURB set was only required to contain a *weak* best response to every strategy that it includes, in which case every stable convention would be CURB. However, the converse would still fail to hold, because any strategy outside of such a modified CURB set need only be matched by some CURB strategy against every CURB strategy profile, where the matching strategy need not be constant independent of the counterparts’ play (as is the case for a stable convention). For this reason, the two concepts only coincide for single-valued solutions, i.e. a singleton CURB set is equivalent to a singleton stable convention, which in turn coincides with a strict Nash Equilibrium.

Indeed, one of the most prominent interpretations of the Nash solution is that of a ‘socially stable convention’, emerging over time as agents adjust their behaviour to that of others observed in the past. See Schotter [2008] for an overview of the related literature, and Young [1993] for a seminal study of the stochastic processes involved. The current paper extends the underlying static notion, based on the premise that, in reality, institutions / conventions do not determine human activities exactly, but merely provide a framework allowing for a certain flexibility of behaviour within the imposed constraints. In fact, a direct parallel can be drawn between the present concept and one of the earliest and most profound formalizations of standards of behaviour.

4.2.2 Relationship to vN-M stable sets

Definition 8 can be linked to the vN-M stable set concept, by expressing it using the general systems framework by Luo [2001]. To see this, consider the following instance of

a general system:

$$\left(S, \{ \succ^A \}_{A \subseteq S} \right), \quad S \equiv \times_{j \in N} S_j \quad (4.1)$$

and let:

$$A_{-i} \equiv \{ s_{-i} \in S_{-i} \mid \exists s_i \in S_i : (s_i, s_{-i}) \in A \}$$

where \succ^A is a conditional dominance relation on S such that for any $s^1, s^2 \in S$ and $A \subseteq S$ we have $s^1 \succ^A s^2$ if and only if for some $i \in N$ it is true that for all $a_{-i} \in A_{-i}$ we have $U_i(s_i^1, a_{-i}) \geq U_i(s_i^2, a_{-i})$, and the inequality is strict for some a_{-i} .

The general stable set $S' \subseteq S$ of system (4.1) is then defined as the vN-M stable set of an abstract game $(S, \succ^{S'})$, i.e. one in which the unconditional dominance relation coincides with $\succ^{S'}$. That is to say, S' must satisfy:

1. [internal stability] $\nexists s, s' \in S' : s \succ^{S'} s'$
2. [external stability] $\forall s \in S \setminus S' \exists s' \in S' : s' \succ^{S'} s$

Firstly, observe that S' must be in Cartesian product form, i.e.: $S' \equiv \times_{j \in N} S'_j$. Otherwise, any strategy profile not in S' but consisting of strategies that are each part of some strategy profile in S' could only be dominated by an element of S' if the internal stability requirement was violated.

Clearly, the internal stability requirement is then equivalent to the ‘if’ part of Definition 8, which states that one conventional strategy cannot weakly dominate another. Similarly, any S' satisfying the ‘only if’ part of the Definition must be externally stable, because any strategy profile $s \notin S'$ includes a strategy s_i weakly dominated by some $s'_i \in S'_i$, so that $s' \succ^{S'} s$ for any strategy profile $s' \in S'$ which includes s'_i . Conversely, for any strategy $s_i \notin S'_i$, where S' is generally stable, consider a strategy profile

$$s \equiv \{s_j\}_{j \in N}, \quad \text{where } s_j = s_i \text{ for } i = j \text{ and otherwise } s_j \in S'_j.$$

There must exist a $s' \in S'$ such that $s' \succ^{S'} s$ and, as ensured by the internal stability of S' , such that $s'_j \in S'_j$ for $i \neq j$. Consequently s_i must be weakly dominated by some $s'_i \in S'_i$, satisfying the ‘only if’ part of Definition 8.

It follows that any stable convention is a vN-M stable set for an unconditional dominance relation on S specified so as to coincide with $\succ^{S'}$. The convenient terms of internal and external stability will also, henceforth, be used in reference to the respective parts of Definition 8.

4.3 Application: Cournot vs. Bertrand

4.3.1 The model

Consider a simultaneous-move variant of the Moreno and Ubeda [2006] modification of Kreps and Scheinkman [1983]. In particular, suppose there are $n \geq 2$ firms in the market, characterized by an inverse demand function $P(\cdot)$ that is twice continuously differentiable, strictly decreasing and concave so long as it remains positive, and equal to zero thereafter. Unlike in the above models, firms are assumed to make both quantity and pricing decisions simultaneously, so that the quantity choice need not be interpreted as a capacity precommitment. Thus, a firm's strategy is a non-negative quantity-price pair, i.e.:

$$S_i \equiv \{(q_i, p_i) \mid q_i, p_i \geq 0\}$$

In addition, production technology is no longer required to be the same for all firms, where the individual cost functions $C_i(\cdot)$ are still twice continuously differentiable and convex, while satisfying:

$$\forall i \in N : 0 < C'_i(0) < P(0) \wedge C_i(0) = 0$$

An important feature is that the firms' unit prices p_i are the minimum at which they declare themselves willing to sell their corresponding outputs q_i . This gives rise to an aggregate supply correspondence $S(p)$, defined for a profile of quantity-price strategies and a market price p as follows:

$$S(p) = \left[\sum_{j \in \{i \in N \mid p_i < p\}} q_j, \sum_{j \in \{i \in N \mid p_i \leq p\}} q_j \right]$$

Although $S(p)$ may not be single-valued, the equilibrium price p^* is uniquely determined by the market-clearing condition:

$$P^{-1}(p^*) \in S(p^*)$$

As opposed to the Kreps and Scheinkman [1983] specification, output is sold exclusively at the market equilibrium price. In particular, when the quantity demanded $P^{-1}(p^*)$ is insufficient to match the entire produce of firms with reservation prices not greater than p^* , only firms which set prices strictly below p^* sell their entire output. Any residual demand is distributed among the firms characterized by $p_i = p^*$ according to a tie-breaking rule, which is left unspecified, as it turns out that it does not influence the results in any way. More precisely, let x_i denote the quantity of output sold by firm i . We then have $x_i = 0$

for $p_i > p^*$, $x_i = q_i$ for $p_i < p^*$, and finally, for $p_i = p^*$, the exact value of x_i depends on the tie-breaking specification, but lies within the following interval:

$$\left[\max \left\{ 0, P^{-1}(p^*) - \sum_{j \in \{l \in N \setminus \{i\} \mid p_l \leq p^*\}} q_j \right\}, \min \left\{ q_i, P^{-1}(p^*) - \sum_{j \in \{l \in N \mid p_l < p^*\}} q_j \right\} \right]$$

Although this entails an implicit rationing rule (firms with $p_i < p^*$ sell first), it is the only appropriate rule in the present circumstances and by no means arbitrary. Indeed, one cannot imagine a situation in which output offered below the prevailing market price remains unsold.

As argued by Moreno and Ubeda [2006], this type of pricing specification describes more realistically situations in which most of the trading occurs at the market clearing price, where the authors point to utility industries as an example. In my view, confining the game to a single stage makes the model even more widely applicable to modern-day trading platforms. For instance, in eBay, as well as other online auction sites, sellers can offer their products at a minimum/starting price, specified together with other auction parameters. Unsurprisingly, the final selling prices of close substitutes are then equal or marginally differentiated. In fact, a similar mechanism is intrinsic to many other markets, from the opening/closing stock market auctions to online betting exchanges.

The last issue that should be considered at this point is whether production costs incurred by the firms are ‘sunk’, i.e. incurred at the time of making an offer, based on the values of q_i , or depend only on the actual sales $x_i \leq q_i$. The two respective payoff specifications are:

$$\begin{aligned} \text{[sunk costs]} & : \pi_i = p^* x_i - C_i(q_i) \\ \text{[sales-dependent costs]} & : \pi_i = p^* x_i - C_i(x_i) \end{aligned}$$

Although the above mentioned models adhere to the first alternative (as is natural when interpreting quantity as capacity precommitment), it is an empirical matter which specification is appropriate for simultaneous price-quantity decisions. For instance, some (but not all) types of goods may need to be produced ‘up-front’, e.g. before setting up an auction, so as to enable the seller to deliver them within a reasonable time frame. However, in the latter case it may still be possible to recover a certain fraction of the costs of producing the output that has not been sold. For this reason, the present paper will consider both of the above extreme possibilities, as well as (for constant unit costs) the ‘intermediate’

cases, i.e. the general payoff function that will be used is:

$$\pi_i = p^* x_i - [(1 - \gamma) C_i(x_i) + \gamma C_i(q_i)]$$

where $\gamma \in [0, 1]$ is the proportion of the production cost that is sunk and cannot be recovered for unsold output.

4.3.2 Results

A convenient feature of the current model specification is that by restricting the original strategy sets S_i accordingly, one can reduce the game to a Cournot or a Bertrand model. In particular, a collection $\{S'_i\}_{i \in N}$ of restricted strategy sets $S'_i \subset S_i$ is compatible with the Cournot specification only if for any strategy profile

$$\{(q'_i, p'_i)\}_{i \in N} \in S' \equiv \times_{i \in N} S'_i$$

the total of players' outputs $Q = \sum_{i \in N} q'_i$ is sold at the corresponding demand price $p^* = P(Q)$. This in turn happens only when the demand price $P(Q)$ exceeds any of the players' respective reservation prices p'_i .

Furthermore, it should be the case that each player's strategy set S'_i includes a continuum of quantities up to the optimal monopoly output (larger quantities are strictly dominated in a Cournot game under the present assumptions). More precisely:

Definition 10 Let q_i^m denote the optimal monopoly output of player i , i.e.:

$$q_i^m = \arg \max_{q_i \geq 0} \{q_i P(q_i) - C_i(q_i)\}$$

A collection of restricted strategy sets $\{S'_i\}_{i \in N}$ is **Cournot-compatible** when we have:

$$\forall i \in N \quad \forall q'_i \in [0, q_i^m] \quad \exists p'_i \quad : \quad (q'_i, p'_i) \in S'_i \quad (C1)$$

$$\forall \{(q'_i, p'_i)\}_{i \in N} \in S' \quad : \quad P\left(\sum_{i \in N} q'_i\right) \geq \max_{i \in N} p'_i \quad (C2)$$

Despite the fact that players' strategies are specified in terms of both quantities and prices, satisfying the above definition means that there is a one-to-one mapping between such quantity-price pairs and their corresponding quantities in the classic Cournot game, in the sense that they each yield the same payoffs in their respective games. In other words, a Cournot-compatible restriction of the current game may be thought of as taking the

original Cournot model and ‘labelling’ every quantity strategy with an additional number (price) without any consequence for the payoffs. Whether or not these labels are present, we are then dealing with the same game.

We now have the apparatus required to investigate the question of whether the Cournot model is a justifiable restriction of a more general price-quantity competition framework. More specifically, I will now establish the necessary and sufficient conditions for the existence of a Cournot-compatible restriction of the model defined in Section 4.3.1, which is also stable in the sense of Definition 8.

Proposition 11 *For $\gamma = 1$ (completely sunk costs), Cournot competition constitutes a stable convention if and only if the following two conditions hold for every player $i \in N$:*

$$P\left(\sum_{j \neq i} q_j^m\right) - C'_i(0) \leq 0 \tag{IS1}$$

$$P\left(\sum_j q_j^m\right) + q_i^m P'\left(\sum_j q_j^m\right) \geq 0 \tag{ES1}$$

For $\gamma = 0$ (completely sales-dependent costs), Cournot competition is never stable.

Proof. See the Appendix. ■

It is easy to see that condition IS1 is both necessary and sufficient for the existence of an internally stable (satisfying the ‘if’ part of Definition 8) Cournot-compatible convention. On the one hand, if it did not hold for some player i , then producing nothing would be dominated by a sufficiently small (and hence also in S'_i) quantity of output. On the other hand, if it does hold for every player, then a convention comprising exclusively quantities below the monopoly optimal ones is internally stable, because no such level of output dominates another in a Cournot game.

Condition ES1 is similarly linked to the external stability (‘only if’) requirement, albeit in a somewhat more complicated way. Intuitively, the potential appeal of Cournot competition to the participating firms is that it ensures that one’s entire output is always sold. This is particularly important when production costs are sunk and the firms tend to sell relatively small quantities of output at relatively high prices. To see why, observe that it is then better to sell everything at a lower (but still relatively high) price, than to leave some produce unsold, but succeed in charging more for the remainder.

Such a low ratio of quantities to prices can, in turn, be associated with an inefficient production technology and a small number of firms in the industry. In such case, even the optimal monopoly outputs are low (individually and in total), as reflected by condition ES1 (recall that the inverse demand $P(\cdot)$ is concave, making the inequality true for sufficiently

small quantities q_i^m). In particular, the condition states that the marginal revenue of each firm is non-negative even when everyone produces their largest rationalizable output. This happens when the marginal costs of production rise quickly, as the firms only want to produce while their monopolistic marginal revenue is at least as large. When the number of players is not too big, then even the substantially smaller ‘oligopolistic’ marginal revenue is still non-negative.

In contrast, when the monopoly outputs are high and there are many firms, Cournot competition is potentially risky, as it entails a possibility of a damaging ‘coordination failure’ scenario, in which a very large aggregate output means the firms sell their substantial produce at a price significantly below their average cost level. Hence, in such case Cournot is no longer a stable convention.

For more insight into the origin of condition ES1, note that it states that when every firm produces q_i^m , the decrease in revenue due to a marginal reduction of sales (for a fixed price) is greater than the decrease in revenue brought about by a marginal increase in the competitors’ production, subject to a price adjustment allowing the increased aggregate output to sell entirely. In other words, a surge in the counterparts’ production is less damaging to players who compromise on their products’ prices, than to ones who would rather maintain the original unit price at the cost of being left with the residual customer demand. Indeed, the fact that this holds when everyone behaves as a monopolist implies the same is true for lower production levels (see Proof).

Proposition 11 establishes that Cournot competition can be a stable convention when costs are sunk, but not when they are sales-dependent. This suggests that, in general, Cournot may be stable when the proportion γ of costs that are impossible to recover for unsold output is sufficiently large. Intuitively, when γ increases it becomes less appealing to sell less at the same price (while losing a proportion γ of the unsold output’s cost), compared to selling the original amount cheaper. Indeed, the next proposition confirms this intuition for constant unit costs of production³.

Proposition 12 *For $\gamma \in [0, 1]$ and $C_i(q) = qc_i$ (constant unit costs), Cournot competition constitutes a stable convention if and only if the following conditions hold for all $i \in N$:*

$$P\left(\sum_{j \neq i} q_j^m\right) - c_i \leq 0 \quad (\text{IS2})$$

$$P\left(\sum_j q_j^m\right) + q_i^m P'\left(\sum_j q_j^m\right) - (1 - \gamma)c_i \geq 0 \quad (\text{ES2})$$

³For general convex costs I was unable to simplify the conditions sufficiently so as to give them a clear economic interpretation.

To further illustrate how production costs and the number of players affect stability of Cournot competition, consider the following ‘textbook’ example of linear demand and constant unit costs.

Example 13 *Suppose that $P(Q) = \max\{\alpha - \beta Q, 0\}$ and for each $i \in N : C_i(q_i) = q_i c$, where $\alpha > c > 0, \beta > 0$, so that $\pi_i = p^* x_i - c_i [(1 - \gamma) x_i + \gamma q_i]$.*

It follows that $\forall i \in N : q_i^m = (\alpha - c) / 2\beta$, so that condition(s) IS2 reduce to $n \geq 3$, while condition(s) ES2 are equivalent to:

$$\frac{c}{\alpha} \geq \frac{n - 1}{n - 1 + 2\gamma}$$

In other words, Cournot competition is a stable convention when the unit cost of production and the proportion of sunk costs are large relative to the number of players and the maximum price the consumers are ever willing to pay.

It is interesting to compare the results in Propositions 11 and 12 to those of Moreno and Ubeda [2006]. They show that, despite the multiplicity of equilibria at the price competition stage, a pure strategy equilibrium always exists, while any pure-strategy equilibrium of the two-stage game yields the Cournot outcome. This improvement on the Kreps and Scheinkman [1983] result (where pure strategy equilibria in prices may fail to exist) is achieved by introducing reservation, rather than exact pricing.

It turns out that reservation prices are equally instrumental in providing an alternative justification for Cournot competition. In fact, the current framework demonstrates that the entire Cournot game, rather than its outcome alone, may constitute a self-sustaining convention within a more general form of price-quantity competition. This is achieved without restructuring the original problem with respect to the timing of moves, and relying on equilibrium prices to follow any given vector of revealed outputs (capacities). Instead, the belief that others are willing to sell their (unknown) outputs relatively cheap is sufficient to make it rational for each firm to set its own reservation price similarly low, provided the conditions of Proposition 11 or 12 hold. This results in a strategic interaction equivalent to Cournot competition, thereby setting the stage for the realisation of the associated, empirically plausible equilibrium outcome.

The results are also, to an extent, analogous to those of Klemperer and Meyer, who found that, in the presence of demand uncertainty, equilibria exhibit Cournot-like supply-function strategies when the number of firms is small and the (increasing) marginal cost curves are steep relative to demand. This in turn can generally be related to the total of monopoly-optimal outputs being relatively small, which is what ensures that Cournot

competition is externally stable in the present context. Interestingly, while Klemperer and Meyer rely on demand uncertainty, the current results are obtained under uncertainty about the actions of the competitors.

In a similar fashion, I will now apply Definition 8 to evaluate the alternative, Bertrand specification of oligopolistic competition. To this end, let p_i^0 denote the monopoly ‘break-even’ price, i.e. the smallest price such that the associated monopoly profit is non-negative. To fully mimic the Bertrand behaviour, it is necessary that each player’s strategy set S'_i now includes at least a continuum of prices from p_i^0 up to the optimal monopoly price $P(q_i^m)$. Furthermore, the output level for each quantity-price strategy in S'_i must be such that the firm is committed to satisfying the entire customer demand at the associated price. Finally, costs must be sales-dependent, so that a player whose price is undercut does not have to pay for producing the unsold output. More precisely:

Definition 14 *A collection of restricted strategy sets $\{S'_i\}_{i \in N}$ is **Bertrand-compatible** when $\gamma = 0$ and for each $i \in N$ we have:*

$$\forall p'_i \in [p_i^0, P(q_i^m)] \exists q'_i : (q'_i, p'_i) \in S'_i \quad (B1)$$

$$\forall (q'_i, p'_i) \in S'_i : q'_i \geq P^{-1}(p'_i) \quad (B2)$$

Note that condition B2 means that trade occurs at a price p^* equal to the minimum of the players’ reservation prices, as the market-clearing condition $P^{-1}(p^*) \in S(p^*)$ is then satisfied (recall the model specification in Section 4.3.1). The result below easily extends to all $\gamma \in [0, 1]$, and is also valid for any tie-breaking rule applied when two or more players charge the same lowest price.

Proposition 15 *Bertrand competition never constitutes a stable convention.*

Proof. See the Appendix. ■

The above Proposition states that no Bertrand-compatible restriction of the original game can satisfy the stability requirements of Definition 8. The reasoning behind this result is fairly simple. If a player produces less than the market demand at the associated reservation price, then this strategy cannot be weakly dominated by a Bertrand price *above* the unconventional one, since the former is better when rivals set their conventional prices in between the two. Furthermore, the same is true for any conventional price *below* the unconventional one, since the latter yields higher profits when the competitors’ Bertrand prices are above the demand price of the unconventional output, at least as long as the last

price lies below the monopoly optimal one. Thus, the external stability criterion is shown to fail through an argument similar to the one which establishes the internal stability of a Bertrand game.

It follows that Bertrand price competition cannot be justified as a convention within the general quantity-price game that is self-sustaining prior to the formation of conjectures about the competitors' pricing behaviour. In other words, strategies that do not entail a commitment to meet the entire demand at the chosen price cannot be ruled out as irrational based solely on the fact that others will select *some* Bertrand-style prices. This suggests that the earlier mentioned drawbacks of the Bertrand framework, such as its indeterminate or paradoxical predictions ('two is enough for competition'), would not apply to markets in which reservation pricing is possible, as the firms would not frame their strategic interaction in the postulated way.

Naturally, this result depends on the way the general price-quantity model is specified. For instance, if the firms were to choose the *exact* prices at which to offer their output, as in the Kreps and Scheinkman [1983] model, then some Bertrand-compatible convention may well be stable. However, in such a case no restriction of the players' strategy sets could be compatible with the Cournot framework, making such pricing specification unsuitable for the comparative analysis of Cournot and Bertrand as *alternative* simplifications of oligopolistic competition.

4.4 Concluding Remarks

In this paper I put forward a formal criterion for the evaluation and comparison of alternative ways in which the players' strategy sets in non-cooperative games can be restricted to model the strategic interaction at hand in a more simplified manner.

The criterion builds on the premise that, since players make conjectures about the counterparts' strategy choice based on their knowledge of the structure of the game, the latter should be determined in the absence of such conjectures. In particular, a restriction of the players' strategy sets is said to constitute a stable convention when expecting others to choose *some* conventional strategies is sufficient to ensure that each player's range of irrational choices coincides with unconventional behaviour. More precisely, a strategy lies outside the convention if and only if it is weakly dominated by some strategy within a stable convention, so long as others act in a conventional manner. In fact, it was shown that any stable convention constitutes a von Neumann-Morgenstern stable set for a dominance relation over strategy profiles specified accordingly.

The postulated criterion was then used to revisit the Cournot vs. Bertrand controversy in the modelling of oligopolistic competition. In particular, a game was considered in which firms simultaneously set quantities of output and minimum prices at which they would be willing to sell it. It turned out that a restriction of the default strategy sets equivalent to a Cournot game is a stable convention when production costs are high relative to the number of firms, and difficult to recover when output is not sold. In other words, the belief that others are willing to sell their outputs relatively cheap can be sufficient to make it rational for every player to do likewise, thereby implementing Cournot competition. Thus, the entire Cournot game, rather than just its equilibrium outcome, can be justified as an approximation of a more general simultaneous-move quantity-price competition, making it possible to argue that the Cournot model may in some cases be ‘right for the right reasons’. In contrast, Bertrand competition was shown never to constitute a stable convention, suggesting that the associated paradoxical outcome does not apply in the studied circumstances.

These results are encouraging for the application of the postulated model evaluation criterion to other modelling problems and controversies in non-cooperative game theory of oligopolistic competition, such as the plausibility of mixed-strategy pricing or the timing of moves in leader / follower games.

Appendix

Proof of Proposition 11. In order to be stable, a Cournot-compatible convention $\{S'_i\}_{i \in N}$ must not include quantities in excess of the monopoly output, which, due to the concavity of profits, are dominated by other conventional strategies, violating internal stability. Thus, Cournot-compatibility in this case means that every $(q'_i, p'_i) \in S'_i$ satisfies:

$$p'_i \leq P\left(q'_i + \sum_{j \neq i} q_j^m\right)$$

(recall part (C2) of Definition 10).

Condition IS1 is then necessary and sufficient for internal stability of S'_i , ensuring that no $q'_i \in [0, q_i^m]$ is dominated by any other output in a Cournot-compatible game.

I will now show that ES1 is necessary for the external stability of $\{S'_i\}_{i \in N}$. To this end, consider an unconventional strategy $\hat{s}_i = (\hat{q}_i, \hat{p}_i) \in S_i \setminus S'_i$ comprising an output $\hat{q}_i = q_i^m$ and a reservation price \hat{p}_i marginally above the demand price associated with all firms producing their monopoly outputs. The only strategy $s'_i \in S'_i$ that could match it entails

$q'_i = \hat{q}_i = q_i^m$, since any other quantity is outperformed if the remaining firms choose to produce nothing. Let:

$$\hat{Q}_{-i} = P^{-1}(\hat{p}_i) - \hat{q}_i$$

So that the unconventional strategy is equaled by s'_i as long as the aggregate output Q_{-i} from other players' conventional strategies is not greater than \hat{Q}_{-i} . For $Q_{-i} > \hat{Q}_{-i}$, the difference between profits resulting from s'_i and those associated with \hat{s}_i is:

$$\Delta(Q_{-i}) = [q_i^m P(q_i^m + Q_{-i}) - C_i(q_i^m)] - \{\hat{p}_i [P^{-1}(\hat{p}_i) - Q_{-i}] - C_i(q_i^m)\}$$

Thus, for $\Delta(Q)$ to be non-negative for all $Q_{-i} > \hat{Q}_{-i}$, it is necessary that:

$$\Delta'(\hat{Q}_{-i}) = q_i^m P'(q_i^m + \hat{Q}_{-i}) + \hat{p}_i = q_i^m P'(P^{-1}(\hat{p}_i)) + \hat{p}_i \geq 0$$

which holds for \hat{p}_i arbitrarily close to $P(\sum_j q_j^m)$ only if:

$$q_i^m P'(P^{-1}(P(\sum_j q_j^m))) + P(\sum_j q_j^m) \geq 0 \Leftrightarrow -q_i^m P'(\sum_j q_j^m) \leq P(\sum_j q_j^m)$$

which amounts to condition ES1.

In addition, it is necessary that production costs are 'sunk'. Otherwise, the difference between s'_i and \hat{s}_i for $Q_{-i} > \hat{Q}_{-i}$ is:

$$\Delta(Q_{-i}) = [q_i^m P(q_i^m + Q_{-i}) - C_i(q_i^m)] - \{\hat{p}_i [P^{-1}(\hat{p}_i) - Q_{-i}] - C_i(P^{-1}(\hat{p}_i) - Q_{-i})\}$$

leading to:

$$\Delta'(\hat{Q}_{-i}) = q_i^m P'(P^{-1}(\hat{p}_i)) + \hat{p}_i - C'_i(q_i^m) \geq 0$$

or, for \hat{p}_i arbitrarily close to $P(\sum_j q_j^m)$:

$$q_i^m P'(\sum_j q_j^m) + P(\sum_j q_j^m) - C'_i(q_i^m) \geq 0$$

Substituting $q_i^m P'(\sum_j q_j^m) + P(\sum_j q_j^m)$ for $C'_i(q_i^m)$, based on the monopoly profit maximizing condition, this becomes:

$$q_i^m [P'(\sum_j q_j^m) - P'(q_i^m)] + [P(\sum_j q_j^m) - P(q_i^m)] \geq 0$$

which is never the case for a concave $P(\cdot)$.

It remains to show that ES1 is not only necessary, but also sufficient to ensure the external stability of *some* Cournot-compatible convention $\{S'_i\}_{i \in N}$. For simplicity, we may restrict attention to one with strategy sets characterized as follows:

$$S'_i \equiv \{(q'_i, p'_i) : q'_i \in [0, q_i^m] \wedge p'_i \leq \tilde{p}\}, \quad \tilde{p} = P\left(\sum_j q_j^m\right)$$

Suppose that:

$$\hat{q}_i = q'_i \in (0, q_i^m] \quad \text{and} \quad \hat{p}_i > P\left(\hat{q}_i + \sum_{j \neq i} q_j^m\right) \geq \tilde{p}$$

Hence, for $Q_{-i} \leq \hat{Q}_{-i}$ both strategies once again yield equal payoffs, while for $Q_{-i} > \hat{Q}_{-i}$ we have:

$$\Delta(Q_{-i}) = [q'_i P(q'_i + Q_{-i})] - [\hat{p}_i \max\{P^{-1}(\hat{p}_i) - Q_{-i}, 0\}]$$

so that for $Q_{-i} \in (\hat{Q}_{-i}, P^{-1}(\hat{p}_i))$:

$$\Delta'(Q_{-i}) = q'_i P'(q'_i + Q_{-i}) + \hat{p}_i$$

and, using the concavity of $P(\cdot)$:

$$\frac{\partial^2 \Delta}{\partial Q_{-i} \partial q'_i} = P'(q'_i + Q_{-i}) + q'_i P''(q'_i + Q_{-i}) < 0, \quad \frac{\partial^2 \Delta}{\partial Q_{-i}^2} = q'_i P''(q'_i + Q_{-i}) < 0$$

which implies that:

$$\Delta'(Q_{-i}) > q_i^m P'\left(\sum_j q_j^m\right) + \tilde{p} \geq 0$$

where the last inequality follows from condition ES1. Thus, for $Q_{-i} \in (\hat{Q}_{-i}, P^{-1}(\hat{p}_i))$ any increase in Q_{-i} will improve the payoff from using s'_i relative to that of \hat{s}_i , so that the former will outperform the latter. Trivially, this will continue to hold for $Q_{-i} \geq P^{-1}(\hat{p}_i)$, since the revenue generated by the unconventional strategy is then equal to zero, while both strategies entail the same costs.

Consider now the case of $\hat{q}_i > q_i^m = q'_i$ and $\hat{p}_i > \tilde{p}$. For $Q_{-i} \leq \hat{Q}_{-i}$ all firms' entire outputs are sold, so q_i^m is better than \hat{q}_i . For $Q_{-i} \in (\hat{Q}_{-i}, P^{-1}(\hat{p}_i))$, similarly to the previous case:

$$\Delta'(Q_{-i}) = q_i^m P'(q_i^m + Q_{-i}) + \hat{p}_i > q_i^m P'\left(\sum_j q_j^m\right) + \tilde{p} \geq 0$$

and for $Q_{-i} \geq P^{-1}(\hat{p}_i)$ the revenue generated by \hat{s}_i is equal to zero, while the associated costs are higher than $C_i(q_i^m)$.

Finally, suppose $\hat{q}_i > q_i^m = q'_i$ and $\hat{p}_i \leq \tilde{p}$. For $Q_{-i} \leq P^{-1}(\tilde{p}) - \hat{q}_i$ the entire outputs are

sold and s'_i outperforms \hat{s}_i . For $Q_{-i} > P^{-1}(\tilde{p}) - \hat{q}_i$ the conventional strategy gives a profit of $q_i^m P(q_i^m + Q_{-i}) - C_i(q_i^m)$, while the unconventional one yields at most $\tilde{p}\hat{q}_i - C_i(\hat{q}_i)$, which is the case when other players' conventional prices are all at their maximum of \tilde{p} and either $\hat{p}_i < \tilde{p}$ or $\hat{p}_i = \tilde{p}$ and the tie-breaking rule is such that it allocates the maximum possible demand to player i . In other words, the existing output q_i^m is sold at a price of at most \tilde{p} , i.e. not greater than $P(q_i^m + Q_{-i})$, which, for the additional output $\hat{q}_i - q_i^m$, is less than the average cost of its production. This is because condition IS1 implies:

$$\tilde{p} = P\left(\sum_j q_j^m\right) < P\left(\sum_{j \neq i} q_j^m\right) \leq C'_i(0)$$

and $C_i(\cdot)$ is convex. Consequently, \hat{s}_i is, again, outperformed by s'_i , which completes the proof. ■

Proof of Proposition 12. Based on the above proof of Proposition 11, it is clear that condition IS2 and the exclusion of quantities in excess of the monopoly output, are both necessary and sufficient for internal stability. Analogously, I will now show that ES2 is necessary for the external stability of $\{S'_i\}_{i \in N}$, by considering the case of $\hat{q}_i = q'_i = q_i^m$ and \hat{p}_i marginally above $\tilde{p} = P\left(\sum_j q_j^m\right)$. For $Q_{-i} > \hat{Q}_{-i} = P^{-1}(\hat{p}_i) - \hat{q}_i$, the difference between profits resulting from s'_i and those associated with \hat{s}_i is:

$$\begin{aligned} \Delta(Q_{-i}) &= [q_i^m P(q_i^m + Q_{-i}) - c_i q_i^m] - \\ &\quad - \{\hat{p}_i [P^{-1}(\hat{p}_i) - Q_{-i}] - [(1 - \gamma) c_i (P^{-1}(\hat{p}_i) - Q_{-i}) + \gamma c_i q_i^m]\} \end{aligned}$$

Thus, for $\Delta(Q)$ to be non-negative for all $Q_{-i} > \hat{Q}_{-i}$, it is necessary that:

$$\Delta'(\hat{Q}_{-i}) = q_i^m P' \left(q_i^m + \hat{Q}_{-i} \right) + \hat{p}_i - (1 - \gamma) c_i \geq 0$$

which is the same as:

$$q_i^m P' \left(P^{-1}(\hat{p}_i) \right) + \hat{p}_i \geq (1 - \gamma) c_i$$

This holds for \hat{p}_i arbitrarily close to $P\left(\sum_j q_j^m\right)$ only if:

$$\begin{aligned} q_i^m P' \left(P^{-1} \left(P \left(\sum_j q_j^m \right) \right) \right) + P \left(\sum_j q_j^m \right) &\geq (1 - \gamma) c_i \\ \Leftrightarrow q_i^m P' \left(\sum_j q_j^m \right) + P \left(\sum_j q_j^m \right) &\geq (1 - \gamma) c_i \end{aligned}$$

which amounts to condition ES2.

I will now show that ES2 is not only necessary, but also sufficient to ensure the external

stability of a Cournot-compatible convention as follows:

$$S'_i \equiv \{(q'_i, p'_i) : q'_i \in [0, q_i^m] \wedge p'_i \leq \tilde{p}\}$$

Suppose that:

$$\hat{q}_i = q'_i \in (0, q_i^m], \hat{p}_i > P\left(\hat{q}_i + \sum_{j \neq i} q_j^m\right) \geq \tilde{p}$$

It follows that for $Q_{-i} \leq \hat{Q}_{-i}$ both strategies yield equal payoffs, while for $Q_{-i} > \hat{Q}_{-i}$:

$$\begin{aligned} \Delta(Q_{-i}) &= [q'_i P(q'_i + Q_{-i}) - c_i q'_i] - \hat{p}_i \max\{P^{-1}(\hat{p}_i) - Q_{-i}, 0\} + \\ &\quad + [(1 - \gamma) c_i \max\{P^{-1}(\hat{p}_i) - Q_{-i}, 0\} + \gamma c_i q'_i] \end{aligned}$$

so that for $Q_{-i} \in \left(\hat{Q}_{-i}, P^{-1}(\hat{p}_i)\right)$:

$$\Delta'(Q_{-i}) = q'_i P'(q'_i + Q_{-i}) + \hat{p}_i - (1 - \gamma) c_i$$

and, using the concavity of $P(\cdot)$:

$$\frac{\partial^2 \Delta}{\partial Q_{-i} \partial q'_i} = P'(q'_i + Q_{-i}) + q'_i P''(q'_i + Q_{-i}) < 0, \quad \frac{\partial^2 \Delta}{\partial Q_{-i}^2} = q'_i P''(q'_i + Q_{-i}) < 0$$

which leads to:

$$\Delta'(Q_{-i}) > q_i^m P'\left(\sum_j q_j^m\right) + P\left(\sum_j q_j^m\right) - (1 - \gamma) c_i \geq 0$$

Thus, for $Q_{-i} \in \left(\hat{Q}_{-i}, P^{-1}(\hat{p}_i)\right)$ any increase in Q_{-i} will improve the payoff from using s'_i relative to that of \hat{s}_i , so that the former will outperform the latter. For $Q_{-i} \geq P^{-1}(\hat{p}_i)$ using \hat{s}_i leads to selling no output, so that we have:

$$\begin{aligned} \Delta(Q_{-i}) &= [q'_i P(q'_i + Q_{-i}) - c_i q'_i] + \gamma c_i q'_i = \\ &= q'_i [P(q'_i + Q_{-i}) - (1 - \gamma) c_i] \geq q'_i \left[q_i^m P'\left(\sum_j q_j^m\right) + P\left(\sum_j q_j^m\right) - (1 - \gamma) c_i \right] \geq 0 \end{aligned}$$

where the last inequality again follows from condition ES2.

Consider now the case of $\hat{q}_i > q_i^m = q'_i$ and $\hat{p}_i > \tilde{p}$. For $Q_{-i} \leq \hat{Q}_{-i}$ all firms' entire outputs are sold, so q_i^m is better than \hat{q}_i . For $Q_{-i} \in \left(\hat{Q}_{-i}, P^{-1}(\hat{p}_i)\right)$ we have:

$$\Delta'(Q_{-i}) = q_i^m P'(q_i^m + Q_{-i}) + \hat{p}_i - (1 - \gamma) c_i > q_i^m P'\left(\sum_j q_j^m\right) + P\left(\sum_j q_j^m\right) - (1 - \gamma) c_i \geq 0$$

and for $Q_{-i} \geq P^{-1}(\hat{p}_i)$ the revenue generated by \hat{s}_i is equal to zero, so that:

$$\begin{aligned} \Delta(Q_{-i}) &= q_i^m [P(q_i^m + Q_{-i}) - (1 - \gamma) c_i] \geq \\ &\geq q_i^m \left[q_i^m P' \left(\sum_j q_j^m \right) + P \left(\sum_j q_j^m \right) - (1 - \gamma) c_i \right] \geq 0 \end{aligned}$$

Finally, suppose $\hat{q}_i > q_i^m = q'_i$ and $\hat{p}_i \leq \tilde{p}$. For $Q_{-i} \leq P^{-1}(\tilde{p}) - \hat{q}_i$ the entire outputs are sold and s'_i outperforms \hat{s}_i . For $Q_{-i} > P^{-1}(\tilde{p}) - \hat{q}_i$ the conventional strategy gives a profit of $q_i^m [P(q_i^m + Q_{-i}) - c_i]$, while the unconventional one yields at most $\hat{q}_i (\tilde{p} - c_i)$, independent of the tie-breaking rule. In other words, the existing output q_i^m is sold at a price of at most \tilde{p} , i.e. not greater than $P(q_i^m + Q_{-i})$, which, for the additional output $\hat{q}_i - q_i^m$, is less than c_i . This is because condition IS2 implies:

$$\tilde{p} = P \left(\sum_j q_j^m \right) < P \left(\sum_{j \neq i} q_j^m \right) \leq c_i$$

Hence, \hat{s}_i is, again, outperformed by s'_i , which completes the proof. ■

Proof of Proposition 15. Consider any player i such that $q_i^m = \max_{j \in N} q_j^m$ and an unconventional strategy \hat{s}_i characterized by $\hat{p}_i = 0$ and $\hat{q}_i = q_i^m + \varepsilon_0, \varepsilon_0 > 0$. Suppose the minimum of the remaining players' conventional (Bertrand-compatible) prices is $\tilde{p} = P(q_i^m + \varepsilon_0)$. Then the profit of firm i resulting from \hat{s}_i is $\hat{\pi}_0 = \tilde{p}\hat{q}_i - C_i(\hat{q}_i)$. By the concavity of monopoly payoffs, any price below \tilde{p} yields a profit below $\hat{\pi}_0$, where the latter is positive for ε_0 sufficiently small.

Thus, the only conventional strategy that could match \hat{s}_i in this situation is $s'_i = (q'_i, p'_i)$ such that $q'_i = \hat{q}_i$ and $p'_i = \tilde{p}$. Furthermore, this occurs only if the price-tie is broken by allocating the entire demand to i (for instance, as the most cost-efficient firm, based on a particular 'winner takes all' sharing rule). Suppose this is indeed the case and consider an alternative situation, in which the minimum of the remaining players' conventional prices is $\bar{p} = \tilde{p} - \varepsilon_1$. In this case, s'_i gives zero profit at best (assuming costs are not sunk), while \hat{s}_i yields $\hat{\pi}_1 = \bar{p}\hat{q}_i - C_i(\hat{q}_i)$. Since $\hat{\pi}_1$ can be arbitrarily close to the optimal monopoly profit, \hat{s}_i will outperform s'_i for ε_0 and ε_1 sufficiently small.

Finally, note that if it was not possible for some player $j \neq i$ to price-undercut $P(q_i^m)$, then internal stability would require that $S'_i = (q_i^m, P(q_i^m))$, i.e. no price competition à la Bertrand would be possible due to the monopolistic position of firm i . ■

This thesis reconsidered some of the most widely debated problems in game theoretic modelling of oligopolistic competition, and showed how they may be addressed in the presence of uncertainty.

In the first part of the study, I utilized a variant of the classic Hotelling model in which the distribution of consumer tastes is subject to a random shift between the two stages of the game. As demonstrated in Chapter 2, demand uncertainty overturns the ‘maximum differentiation’ principle, based on making the products less similar so as to avoid head-to-head price competition. The presence of uncertainty weakens this incentive, because one firm or the other is always disadvantaged by the random shift in consumer tastes, and must therefore price aggressively in order to make up for it, regardless of whether the products are similar or not. In fact, in the limit as uncertainty becomes large, the outcome of the mill-pricing game under consideration converges to the socially-optimal one, traditionally associated with spatial price discrimination. Furthermore, this holds for a class of linear-quadratic transport cost / disutility functions, rather than the quadratic formulation alone.

The role of demand uncertainty in making the market more competitive may extend to situations in which the firms are unaware of the exact distribution of random factors affecting the consumer preferences. In particular, in Chapter 3 I showed that in such case uncertainty still reduces the equilibrium level of product differentiation when the players are sufficiently pessimistic in the Arrow / Hurwicz sense, i.e. they each assign a large enough weight to the lowest-profit scenario, in which the shift in tastes strongly

favours the competitor. In addition, a unilateral increase of a player's pessimism works as a strategic deterrent, and makes that player better off regardless of the resolution of the random factors. This generates a Prisoner's Dilemma situation and suggests that strategic commitment, elimination of underperforming firms or taxation could all lead to the firms being exclusively focused on the worst-case scenario, in which consumer tastes shift towards the competitor to the highest possible extent. As a way of insuring against this possibility, the duopolists then locate relatively close together, even though this compels them to price aggressively.

The results are in contrast with the existing literature, which generally associates demand uncertainty with less competition and sub-optimal product variety. However, they could explain some of the observed industrial behaviour. For instance, in situations where consumer behaviour is subject to frequently changing trends, e.g. in fashion or 'new technology' industries, we usually observe substantial product variety. In contrast, when the consumer preferences are fairly constant, there are often only few, possibly extreme varieties on offer (e.g. low vs. high sugar food products). Finally, as discussed in Chapter 3, the postulated role of the firms' pessimistic approach is in line with the tendencies observed in the sports betting industry.

The results could also have interesting policy implications. In Section 2.2, I presented a possible interpretation of the applied uncertainty specification - as a result of the consumers observing a random signal, suggesting what type of product is best and making them reflect on their initial views. In the light of the current study, such recommendations (e.g. regarding the daily intake of sugar) could be provided by the government in an attempt to stimulate competition in the market and improve social welfare. Interestingly, such guidelines could improve product variety and reduce prices regardless of whether they are actually beneficial to the consumers (e.g. whether or not the provided dietary recommendations are truly best for one's health). What matters is that the potential buyers believe this to be the case and the recommendations are decided upon in a way independent of and obscure to the industry.

The model could also offer insights into the regulation of market research practices. As uncertainty, in the form utilized by the current framework, intensifies, the firms know less about the consumer tastes at the time of product design relative to their knowledge at the pricing stage. In recent years, questions have been raised over the companies' growing potential to gather information about consumer preferences via online services, like social-networking sites. A common argument invoked to justify such practices 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 current study

indicates that while firms will take advantage of the additional information in pursuit of their own strategic objectives, this may be inconsistent with the benefit of the consumers. On the contrary, it is when least is known about consumer preferences that sellers offer products to match those preferences as closely as possible.

Finally, an interesting avenue for future research, sketched in Chapter 3, is the effect of corporate taxation on the firms' location decisions. Introducing a tax-free quota would reduce the importance, for a firm's expected payoff, of those realisations of the uncertainty in which the consumer demand shifts sufficiently towards that firm, and away from its competitor, for the firm's profits to exceed the taxation threshold. Conversely, the outcomes of the uncertainty in which the demand shifts in the rival's direction would become more important, inducing the firm to relocate towards the counterpart, thereby reducing the equilibrium product differentiation and the resulting prices. However, the precise mechanism and comparative statics of this process should be investigated in a separate model.

The second part of the thesis considered situations in which two or more firms are selling identical products, while faced with a downwards-sloping (rather than perfectly inelastic) market demand. In particular, Chapter 4 reconsidered an unresolved discrepancy between the two most prominent models of oligopolistic competition - the Cournot model relying on quantity choice and the Bertrand price-setting framework. To this end, I proposed a model evaluation criterion, based on the von Neumann - Morgenstern stable set concept and the idea that the structure of a strategic interaction must be determined prior to the formation of conjectures about the counterparts' play, i.e. under a form of strategic uncertainty. More precisely, a restriction of the players' strategy sets is said to constitute a 'stable convention' when it is made up of all those strategies that are not irrational based on the belief that others will adhere to the convention and without any knowledge regarding their *exact* behaviour.

The above criterion was applied to a game in which players simultaneously decide on the quantity of a homogeneous good they each wish to produce and the minimum price at which they are willing to sell it. Crucially, the set of quantity-price pairs available to the players can be restricted so as to make the strategic interaction at hand equivalent to either a Cournot or a Bertrand game. When the production costs are 'sunk' upon offering the output for sale, it turned out that a Cournot-equivalent restriction of this general quantity-price game is a stable convention if and only if the corresponding Cournot game satisfies two conditions. Firstly, every quantity below the monopoly-optimal one must be rationalizable. Secondly, the marginal revenue of each firm must not be negative even when everyone produces their monopoly outputs, which in turn may be associated with

large production costs, as this ensures that the said monopoly quantities are low and the corresponding marginal revenue is still high.

The conditions become more stringent when players are able to recover a larger proportion of the production cost of any unsold output. Hence, one could say that, overall, Cournot competition is a stable convention within a more general quantity-price game when the irredeemable cost of having produced the output that is not sold (and hence also any output in general) is large. This is very intuitive, as the distinct appeal of Cournot competition is that it ensures that the entire outputs are always sold. The contribution of Chapter 4 of the thesis was to provide formal foundations for such commonly held beliefs and intuitive justifications of the Cournot framework.

In comparison, I also demonstrated that a Bertrand-equivalent restriction of the price-quantity game under consideration is never a stable convention. This suggests that the paradoxically competitive outcome of this competition mode does not apply to reservation pricing, since in such case the firms would not frame their strategic interaction in the Bertrand manner. Under the studied circumstances, and when the derived Cournot stability conditions are met, the comparative statics results of quantity competition are the ones that seem not only ‘right’ (in the sense of being empirically plausible), but also - for the *right* reasons.

The model evaluation criterion put forward here offers substantial possibilities for future research. Firstly, one could seek an empirical verification of the obtained results. This could entail a study of the data obtained from online auction sites, where sellers decide on the quantities of output to offer and the associated reservation prices, while perfect substitutes are traded at a single market equilibrium price. The hypothesis to scrutinize would be whether the trade in products satisfying the derived Cournot stability conditions is indeed consistent with this form of competition. For better control over the testing process, it might also be prudent to design and run an experimental study along the same lines.

An altogether independent line of investigation would be to apply the proposed concept of a stable convention to other modelling problems within the area of industrial organization. For instance, one could address the controversies associated with the plausibility of price randomization, by investigating the conditions that would make restricting the players’ strategy sets to degenerate pure-strategy prices a stable convention within a more general mixed-strategy pricing game. Another possibility would be to formalize the determinants of the timing of moves in leader / follower games, which in reality is subject to players’ choice, rather than being exogenously imposed on them. Such applications of the postulated model evaluation criterion offer intriguing possibilities for future investigations.

Bibliography

- Acemoglu, D., K. Bimpikis, and A. Ozdaglar (2009): “Price and capacity competition,” *Games and Economic Behavior*, 66, 1–26.
- Aghion, P., M. P. Espinosa, and B. Jullien (1993): “Dynamic duopoly with learning through market experimentation.” *Economic Theory*, 3, 517–537.
- Anderson, S. P. (1988): “Equilibrium existence in the linear model of spatial competition,” *Economica*, 55, 479–491.
- Anderson, S. P., J. K. Goeree, and R. Ramer (1997): “Location, location, location,” *Journal of Economic Theory*, 77, 102–127.
- Anderson, S. P., A. d. Palma, and Y. Nesterov (1995): “Oligopolistic competition and the optimal provision of products,” *Econometrica*, 63, pp. 1281–1301.
- Ansari, A., N. Economides, and J. Steckel (1998): “The max-min-min principle of product differentiation,” *Journal of Regional Science*, 38, 207–230.
- 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.
- Basu, K. and J. W. Weibull (1991): “Strategy subsets closed under rational behavior,” *Economics Letters*, 36, 141–146.

- Baye, M. R. and J. Morgan (2002): “Winner-take-all price competition,” *Economic Theory*, 19, 271–282.
- Bernhardt, D., J. Duggan, and F. Squintani (2007): “Electoral competition with privately-informed candidates,” *Games and Economic Behavior*, 58, 1 – 29.
- Bernheim, B. D. (1984): “Rationalizable strategic behavior,” *Econometrica*, 52, 1007–1028.
- Bernheim, B. D., B. Peleg, and M. D. Whinston (1987): “Coalition-proof nash equilibria i. concepts,” *Journal of Economic Theory*, 42, 1 – 12.
- Bertrand, J. L. F. (1883): “Theorie mathematique de la richesse sociale,” *Journal des Savants*, 499–508.
- 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.
- Bhaskar, V. and T. To (2003): “Oligopsony and the distribution of wages,” *European Economic Review*, 47, 371 – 399.
- Bhaskar, V. and T. To (2004): “Is perfect price discrimination really efficient? an analysis of free entry,” *The RAND Journal of Economics*, 35, 762–776.
- Borgers, T. (1992): “Iterated elimination of dominated strategies in a bertrand-edgeworth model,” *The Review of Economic Studies*, 59, 163–176.
- Brenner, S. (2005): “Hotelling games with three, four, and more players,” *Journal of Regional Science*, 45, 851–864.
- Bresnahan, T. F. (1981): “Duopoly models with consistent conjectures,” *The American Economic Review*, 71, pp. 934–945.
- Bresnahan, T. F. (1989): *Chapter 17 Empirical studies of industries with market power*, *Handbook of Industrial Organization*, volume 2, Elsevier.
- Caplin, A. and B. Nalebuff (1991): “Aggregation and imperfect competition: On the existence of equilibrium,” *Econometrica*, 59, pp. 25–59.
- Casado-Izaga, J. F. (2000): “Location decisions: The role of uncertainty about consumer tastes,” *Journal of Economics*, 71, 31–46.

- Cournot, A. (1838): *Recherches sur les principes mathematiques de la theorie des richesses*, Paris.
- Dasgupta, P. and E. Maskin (1986a): “The existence of equilibrium in discontinuous economic games, i: Theory,” *The Review of Economic Studies*, 53, 1–26.
- Dasgupta, P. and E. Maskin (1986b): “The existence of equilibrium in discontinuous economic games, ii: Applications,” *The Review of Economic Studies*, 53, 27–41.
- d’Aspremont, C., J. J. Gabszewicz, and J. F. Thisse (1979): “On hotelling’s “stability in competition”,” *Econometrica*, 47, 1145–1150.
- Dastidar, K. G. (1995): “On the existence of pure strategy bertrand equilibrium,” *Economic Theory*, 5, 19–32.
- Davidson, C. and R. Deneckere (1986): “Long-run competition in capacity, short-run competition in price, and the cournot model,” *The RAND Journal of Economics*, 17, 404–415.
- Delgado, J. (2006): “Coalition-proof supply function equilibria under capacity constraints,” *Economic Theory*, 29, 219–229.
- Delgado, J. and D. Moreno (2004): “Coalition-proof supply function equilibria in oligopoly,” *Journal of Economic Theory*, 114, 231 – 254.
- Dixit, A. (1986): “Comparative statics for oligopoly,” *International Economic Review*, 27, pp. 107–122.
- Dixon, H. D. (1992): “The competitive outcome as the equilibrium in an edgeworthian price-quantity model,” *The Economic Journal*, 102, 301–309.
- Eaton, B. C. and R. G. Lipsey (1975): “The principle of minimum differentiation reconsidered: Some new developments in the theory of spatial competition,” *The Review of Economic Studies*, 42, 27–49.
- Economides, N. (1984): “The principle of minimum differentiation revisited,” *European Economic Review*, 24, 345 – 368.
- Economides, N. (1986): “Minimal and maximal product differentiation in hotelling’s duopoly,” *Economics Letters*, 21, 67–71.
- Economides, N. (1993): “Hotelling’s main street with more than two competitors,” *Journal of Regional Science*, 33, 303–319.

- Edgeworth, F. (1925): *The pure theory of monopoly, Papers relating to political economy*, volume 1.
- Fabra, N., N.-H. von der Fehr, and D. Harbord (2006): “Designing electricity auctions,” *The RAND Journal of Economics*, 37, 23–46.
- Friedman, J. W. (1988): “On the strategic importance of prices versus quantities,” *The RAND Journal of Economics*, 19, 607–622.
- Gabszewicz, J. J. and J.-F. Thisse (1986): “On the nature of competition with differentiated products,” *The Economic Journal*, 96, 160–172.
- Gabszewicz, J. J. and J.-F. Thisse (1992): *Chapter 9: Location, Handbook of Game Theory with Economic Applications*, volume 1, Elsevier.
- Ghirardato, P., F. Maccheroni, and M. Marinacci (2004): “Differentiating ambiguity and ambiguity attitude,” *Journal of Economic Theory*, 118, 133–173.
- Grossman, S. J. (1981): “Nash equilibrium and the industrial organization of markets with large fixed costs,” *Econometrica*, 49, 1149–1172.
- Harrington, J. E. (1989): “A re-evaluation of perfect competition as the solution to the bertrand price game,” *Mathematical Social Sciences*, 17, 315–328.
- Harter, J. F. R. (1997): “Hotelling’s competition with demand location uncertainty,” *International Journal of Industrial Organization*, 15, 327–334.
- Hoernig, S. (2007): “Bertrand games and sharing rules,” *Economic Theory*, 31, 573–585.
- Hotelling, H. (1929): “Stability in competition,” *The Economic Journal*, 39, 41–57.
- Kaas, L. and P. Madden (2010): “Minimum wages and welfare in a hotelling duopsony,” *Economic Theory*, 43, 167–188.
- Klemperer, P. and M. Meyer (1986): “Price competition vs. quantity competition: The role of uncertainty,” *The RAND Journal of Economics*, 17, pp. 618–638.
- Klemperer, P. D. and M. A. Meyer (1989): “Supply function equilibria in oligopoly under uncertainty,” *Econometrica*, 57, 1243–1277.
- Kohlberg, E. and J.-F. Mertens (1986): “On the strategic stability of equilibria,” *Econometrica*, 54, 1003–1037.

- Kreps, D. M. and J. A. Scheinkman (1983): “Quantity precommitment and bertrand competition yield cournot outcomes,” *The Bell Journal of Economics*, 14, 326–337.
- 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.
- Lambertini, L. (1994): “Equilibrium locations in the unconstrained hotelling game,” *Economic notes*, 23, 438–46.
- Lederer, P. J. and J. Hurter, Arthur P. (1986): “Competition of firms: Discriminatory pricing and location,” *Econometrica*, 54, 623–640.
- Luo, X. (2001): “General systems and phi-stable sets. a formal analysis of socioeconomic environments,” *Journal of Mathematical Economics*, 36, 95–109.
- Luo, X. (2009): “On the foundation of stability,” *Economic Theory*, 40, 185–201.
- Madden, P. (1998): “Elastic demand, sunk costs and the kreps-scheinkman extension of the cournot model,” *Economic Theory*, 12, 199–212.
- Marquez, R. (1997): “A note on bertrand competition with asymmetric fixed costs,” *Economics Letters*, 57, 87 – 96.
- Maskin, E. (1986): “The existence of equilibrium with price-setting firms,” *The American Economic Review*, 76, pp. 382–386.
- 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.
- Merel, P. R. and R. J. Sexton (2010): “Kinked-demand equilibria and weak duopoly in the hotelling model of horizontal differentiation,” *The B.E. Journal of Theoretical Economics*, 10: Iss. 1 (Contributions).
- Moreno, D. and L. Ubeda (2006): “Capacity precommitment and price competition yield the cournot outcome,” *Games and Economic Behavior*, 56, 323–332.

- Moscarini, G. and M. Ottaviani (2001): “Price competition for an informed buyer,” *Journal of Economic Theory*, 101, 457 – 493.
- Novshek, W. (1980): “Cournot equilibrium with free entry,” *The Review of Economic Studies*, 47, pp. 473–486.
- Novshek, W. (1985): “On the existence of cournot equilibrium,” *The Review of Economic Studies*, 52, 85–98.
- Osborne, M. J. and C. Pitchik (1987): “Equilibrium in hotelling’s model of spatial competition,” *Econometrica*, 55, 911–922.
- Pearce, D. G. (1984): “Rationalizable strategic behavior and the problem of perfection,” *Econometrica*, 52, 1029–1050.
- Puller, S. L. (2007): “Pricing and firm conduct in california’s deregulated electricity market,” *Review of Economics and Statistics*, 89, 75–87.
- Rubinstein, A. (1991): “Comments on the interpretation of game theory,” *Econometrica*, 59, pp. 909–924.
- Salop, S. C. (1979): “Monopolistic competition with outside goods,” *The Bell Journal of Economics*, 10, pp. 141–156.
- Saporiti, A. and G. Coloma (2010): “Bertrand competition in markets with fixed costs,” *The B.E. Journal of Theoretical Economics*, 10.
- Schotter, A. (2008): *The Economic Theory of Social Institutions*, Cambridge Books, Cambridge University Press.
- Shapiro, C. (1989): *Chapter 6 Theories of oligopoly behavior, Handbook of Industrial Organization*, volume 1, Elsevier.
- Singh, N. and X. Vives (1984): “Price and quantity competition in a differentiated duopoly,” *The RAND Journal of Economics*, 15, 546–554.
- 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.
- Tabuchi, T. (1994): “Two-stage two-dimensional spatial competition between two firms,” *Regional Science and Urban Economics*, 24, 207 – 227.

- Tabuchi, T. and J.-F. Thisse (1995): “Asymmetric equilibria in spatial competition,” *International Journal of Industrial Organization*, 13, 213–227.
- Tirole, J. (1988): *Industrial Organization*, Cambridge: MIT Press.
- Vives, X. (1989): “Cournot and the oligopoly problem,” *European Economic Review*, 33, 503–514.
- Young, H. P. (1993): “The evolution of conventions,” *Econometrica*, 61, 57–84.