Tilburg University

Still searching for Bertrand equilibrium

van Damme, E.E.C.

Published in:
Equilibrium, Markets and Dynamics. Essays in Honour of Claus Weddepohl

Publication date:
2002

Document Version
Publisher's PDF, also known as Version of record

[Link to publication in Tilburg University Research Portal](#)

Citation for published version (APA):
van Damme, E. E. C. (2002). Still searching for Bertrand equilibrium. In C. H. Hommes, R. Ramer, & C. Withagen (Eds.), *Equilibrium, Markets and Dynamics. Essays in Honour of Claus Weddepohl* (pp. 167-182). Springer Verlag.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Still Searching for Bertrand Equilibrium*

Eric van Damme

CentER, Tilburg University

Abstract. Throughout his career, Claus Weddepohl worked on models of price competition. In one of his earliest papers, he stated his motivation for this work: if economists do not succeed in constructing tractable models of price competition that generate rich insights, policy makers will assume that competition among the few will produce competitive outcomes. In this selective survey of Bertrand type models, I address the question whether price competition will produce competitive outcomes, what insights can be derived from these models, and whether a unified theory should be aimed for.

Key words: oligopoly theory, Bertrand competition, non-cooperative games.

1 Introduction

Price competition and the Bertrand model have played an important role in the work of Claus Weddepohl. In one of his first publications (Weddepohl, 1975), he writes (in my (abbreviated) translation of his Dutch):

"The theory of price formation contains three compartments. The first includes the theory of perfect competition, which is supposed to apply when there are many suppliers. The second is the theory of pure monopoly. This theory is simple and old. The third compartment is a mixed bag; here one finds theories that deal with markets with more than one, but not many competitors: oligopoly and monopolistic competition. The many theories which cover this field do not form a unity and they are complicated, some even very complicated. The consequence is that many observers believe that in cases like these, where there are "not too many" but also "not too few" competitors, price formation will more or less be as in the case of perfect competition".

In his paper, Weddepohl then continues by studying a Bertrand model in which firms produce with increasing returns to scale technologies, and he shows that the outcomes are very different from those resulting under perfect competition. In particular, his model admits many equilibria, at different price levels, and with different welfare properties. One can, hence, draw the broad conclusion that it is not wise to assume that markets in which there is competition among the few function as if they were perfectly competitive.

* The title is inspired by one of Claus Weddepohl's contributions "Op zoek naar evenwicht in een Bertrand model".

Weddepohl's article concludes by remarking that many variations of the basic model seem worthy of study and, indeed, later in his career, Weddepohl has returned to this general theme several times; several of his papers discuss variants of Bertrand's model.

Writing a quarter of a century later, I wish to return to this theme, to see what we have learned since then. The motivation for doing so is given by the final sentence from the above quotation, which I interpret as saying that, if economists do not succeed in making simple theories of competition among the few, policy makers will be inclined to assume that these markets will function as if they are perfectly competitive. The Dutch policy experience confirms this assessment of Weddepohl. Since the beginning of the 1990s, the Dutch government has embarked on a program of regulatory reform, the aim being to increase welfare, employment and growth by making markets more competitive. As I have argued at length elsewhere (Van Damme, 2001ab), policy has concentrated on reducing entry barriers, the underlying idea being that, once such barriers had been eliminated, competitive markets would arise spontaneously. The experience, however, has shown this not to be the case; where legal entry barriers were scrapped, strategic entry barriers were erected, so as to avoid competition becoming too intense. Furthermore, firms were also well aware that, in situations of tight oligopoly, it may be better to (tacitly) cooperate rather than compete, and this also contributed to the benefits of competition not being fully realised.

Returning to theory, the good news is that models are simple. The other side of the coin is that there is no unified theory; we just have a collection, a large collection, of models. Each model is reasonably easy to understand, but what may be more complicated is to exercise good judgement about which model is, or which insights are, most relevant in a certain specific case. One may expect that those who know the ins and outs of the various models, are also in the best position to judge which model fits best and which insights really apply. With this in mind, this paper provides an overview of some of the simple models of price competition, and the insights that can be derived from them.

The remainder of the paper is organized as follows. In Section 2, we discuss the basic Bertrand model together with some simple variations, including those that have been studied in the work of Weddepohl. In Section 3, we discuss issues of timing and cooperation in models of price competition, and in Section 4 we briefly turn to behavioral issues: how are price competition games played in the laboratory? While the material in Section 2 is drawn from the literature at large, the discussion in Sections 3 and 4 is mainly based on work that has been done at Tilburg University and in which I have been involved more closely. Section 5 concludes with some general remarks, inspired by the state of the art in this field and by the separation of economics and econometrics in the Netherlands.

2 Variations on Bertrand's Theme

We start with the most simple model of price competition. There are two firms that each can costlessly produce one unit of an indivisible commodity; there is one consumer that is willing to pay up to one money unit for one unit of this commodity. If the consumer buys from firm i for a price p_i , then the utility of firm i is p_i , the utility of the consumer is $1 - p_i$, and the utility of firm j ($j \neq i$) is zero. The game proceeds as follows: simultaneously, each of the firms chooses its price; the consumer observes both prices and then decides to buy from one of the producers at the price quoted there, or not to buy at all.

It is remarkable how rich a set of issues can already be addressed in this setting. First of all, there are issues in pure (game) theory. Let us start by simplifying the game by eliminating the choice of the consumer, assuming that he will always choose what is in his best interest. Hence, if $\min(p_1, p_2) > 1$ the consumer does not buy; otherwise he buys at a firm charging the lowest price. In effect, we are moving from Nash equilibrium to subgame perfect equilibrium. As the reader may verify, we are eliminating some Nash equilibria in this way. (For example, consider a consumer who "wants to buy" at $\frac{1}{2}$; if both firms charge $\frac{1}{2}$, he randomizes equally among them; if exactly one firm charges $\frac{1}{2}$, he buys at that firm; and if none of the firms charges $\frac{1}{2}$, he does not buy.) Hence, one sees that the difference between Nash and subgame perfect equilibria already appears in this very simple model. Already in a simple setting like this one, one can illustrate "incredible threats" and "unreasonable equilibria".

Let us assume that the consumer indeed always behaves optimally, so that we can focus on competition between the firms. A formal definition of the game requires a specification of the strategy sets. Let us first assume that each firm can choose prices from a finite set P . As is well-known, as long as the price grid is sufficiently fine, the game is (weakly) dominance solvable, i.e. by repeatedly eliminating weakly dominated strategies, one can eliminate all prices other than the lowest, strictly positive and feasible price. Hence, one does not need an equilibrium assumption to solve the game. One may discuss several questions related to this observation: Is the procedure of repeatedly eliminating dominated strategies "justified", and under what conditions are the outcomes of such a procedure independent of the order in which one eliminates strategies? Both questions have received ample attention in the recent literature, see for example Brandenburger and Keisler (2001), and Dufwenberg and Stegeman (2001) for recent contributions.

Of course, the choice of a finite grid may be somewhat artificial. Also the case in which P is a continuum, say $P = [0, 1]$, is interesting. In this case, the unique Nash equilibrium (of the reduced game with the consumer eliminated) is $(0, 0)$, but, in a certain sense 0 is the worst strategy for each firm: it is weakly dominated by any other pure strategy in the game. Hence, here we have a paradox and an essential difference between games with a finite strategy set

and games in which strategies can be chosen from a continuum. In the former case, there always exists a Nash equilibrium in undominated strategies, in the latter case apparently not. Given the fact that some Nash equilibria are unreasonable, given the need to refine the Nash equilibrium concept, given perfectness as an acceptable refinement criterion in finite games, and given that perfect equilibria of finite games are undominated, the question arises how to meaningfully define perfect equilibria for continuum games. See Simon and Stinchcombe (1995) for an answer to this latter question.

The economic message that comes from the above example is that price competition is a very competitive situation: with constant marginal cost, already two firms are sufficient to generate the competitive outcome and zero profit. To investigate how general this message is, we now study a couple of variations of the basic model. First of all, we note that, if there are more than two firms that can produce the commodity, we can observe price dispersion in equilibrium: as long as two firms each charge a price of zero, in equilibrium, all of the other firms can randomize in any way they want. (Of course, they won't necessarily make sales.) In the remainder of this Section, we first consider more general supply and demand conditions, whereafter attention is focused on the information that is available to market parties, on issues of divisibility, and on repetition of the game.

It is generally thought that, with constant returns to scale in production, the game is (almost) dominance solvable for rather general demand functions. That this, however, is not true can be seen by considering the case in which the consumer is willing to buy at any price, hence, he will always buy from the cheapest source, and the game played between the two firms is given by $\pi_i(p_1, p_2) = p_i$ if $p_i < p_j$, $\pi_i(p_1, p_2) = 0$ if $p_i > p_j$, and $\pi_i(p_1, p_2) = p_i/2$ if $p_i = p_j$. This game has multiple equilibria, in fact multiple mixed strategy equilibria. To verify this, assume that player i draws a price at random from the distribution $F(\cdot)$. Then player j 's expected payoff resulting from price p is $p(1 - F(p))$. Hence, if $F(p) = 1 - K/p$ for $p \geq K$, then player j 's expected payoff is K for any such price, while, furthermore, the expected payoff is lower for any $p < K$. Hence, we see that for any $K > 0$, we have such a mixed strategy equilibrium. Baye and Morgan (1999) show that a similar construction is possible for more general, downward sloping, demand functions, provided that monopoly profits are unbounded.

In his work, Weddepohl has investigated the relation between the competitive outcome and the Bertrand equilibrium for more general cost functions. Assume the consumer side is represented by a decreasing demand function, $d(p)$. Furthermore, assume that there are n identical firms that all produce with a cost function c . In Weddepohl (1975), the paper that was referred to in the Introduction, Weddepohl analyses the case of increasing returns to scale. In this case c is increasing and concave and the least costly way of serving demand is through monopoly. Weddepohl remarks that, in this case, supply functions and, hence, competitive equilibria, are not well-defined. The

Bertrand price setting game, however, is well-defined. If $n > 1$, there is competition for the market and this eliminates monopoly profits: there is an equilibrium in which one firm is active and charges price p such that $pd(p) = c(d(p))$. The other firms charge a price slightly above p , which ensures that they don't have demand and that the "active firm" cannot profitly deviate. (Note that in case the price space is a continuum, this equilibrium involves the demand side solving a coordination problem -all consumers have to go to the same firm- hence, consumers cannot simply be assigned equally to all firms with the lowest price.) Weddepohl, however, concludes that infinitely many outcomes are possible. It is not that his analysis is incorrect, he analyses a different model. Formally, he considers the following 2-stage model, which is similar to the one from Marschak and Selten (1974). In a first stage, firms decide whether or not to become active and active firms determine their prices. In a second stage, those firms that do not have the lowest price have the opportunity to switch to the lowest price. The consumer then chooses to buy at the lowest price where he equally splits his demand over all firms that do charge this lowest price. Weddepohl (1975) shows that there is a great many equilibria and that the more firms are active, the higher the price is.

In Weddepohl (1991), the paper that motivated the title of the present contribution, Weddepohl considers the case of decreasing returns to scale. Let us assume that there are n firms, that each produce with an increasing and convex cost function $c(\cdot)$ with $c(0) = 0$, and that consumers will distribute evenly across all firms that charge the lowest price. In this case, the competitive equilibrium is well-defined, but each firm charging the competitive price and producing the competitive quantity will typically not be an equilibrium of the price setting game. The reason is simple: if all firms charge the competitive price and produce the competitive quantity, each individual firm has some monopoly power over its own customers and it will have an incentive to exploit that power and to raise its price. As Weddepohl remarks, the equilibrium of the Bertrand game might depend on the fine details of the rationing rule that is used and an equilibrium in pure strategies will typically not exist. (For the special case of capacity constraints, this was already remarked in Edgeworth (1897).) Weddepohl then searches for price vectors that possess unidirectional stability (hence, either unilateral price increases or price decreases are not profitable), but also general existence of such "relaxed equilibria" cannot be established. In related work, Canoy and Weddepohl (1995), a relation "relaxation" is studied, in which firms assume that price reductions are always matched, but price increases not necessarily. We will not discuss these technically complicated aspects here, but confine ourselves to the remark that mixed strategy Nash equilibria can be expected to exist (Dasgupta and Maskin, 1986), even though they may be complicated to characterize or compute explicitly. However, it is straightforward to characterize equilibrium in case firms are not allowed to ration demand, hence,

in case they are forced to serve all demand. (Of course, such an assumption may not be a reasonable one to make.) In this case, if there are m firms that charge the lowest price and if this lowest price is p_m , then each of these firms has profits equal to $\pi_m = p_m q_m - c(q_m)$ where $q_m = d(p_m)/m$, while the others have profit zero. Define p_m^* as that price that would result in zero profits if it is the minimum one in the market and it is charged by exactly m firms. These prices are well-defined under quite general conditions and, as can be easily verified, they are all equilibrium prices. Hence, we have multiple equilibria, see Dastidar (1995).

Another important modification of the basic model concerns the information condition. We have assumed that the consumer is informed about both prices and that firms know this. Imagine instead that the consumer only learns of one of these prices, with both possibilities being equally likely. After having learned p_i , the consumer may decide to buy at that price, or he may, at a cost $\varepsilon > 0$, decide to inquire information about the competitor's price p_j . After the consumer has bought the additional information, he then has to decide whether to buy the product, or not; in this case, if the consumer buys at price p_k his net utility is $1 - p_k - \varepsilon$. In a symmetric, pure strategy Nash equilibrium, the consumer expects the same prices at both locations, hence, he has no reason to switch. Since player i 's best response against p_j is given by $p_i = p_j + \varepsilon$ if $p_j \leq 1 - \varepsilon$, it follows that the unique equilibrium is $(1, 1)$: each firm will charge the monopoly price. The example (originally due to Diamond (1971)) illustrates the importance of market transparency for competition: in a transparent market, the competitive price will result as soon as there are two firms; when the market is not transparent, the price will be the monopoly price, irrespective of the number of firms. Of course, in this stark form, the result may not be too relevant for actual markets. Nevertheless, this simple model may serve as a building block from which essential insights relating to market transparency, and the importance of the Internet for competition may be derived, see Van Damme and Dellaert (2001).

Of course, the situation in which n sellers compete for one client is mathematically equivalent to the situation in which n buyers compete for one object. The latter is the situation analysed in auction theory, hence, the close relation between that theory and the theory of Bertrand competition. Let us focus on the simple auction setting in which one object is available for sale with two buyers computing for it, and each of them assigning value 1 to this object. If the object is indivisible, and the seller organises a sealed-bid, first price auction, we are back in the previous case: each bidder will bid 1. What would happen if the object would be perfectly divisible with buyers assigning marginal value of 1 to each unit? In this case, each bidder i can specify a demand function $q_i(p)$ stating how many units he would like to have in case the seller charges prices p . Assume the seller organises a competition in demand functions, i.e. buyers simultaneously determine their demands and the seller then determines the price that clears the market; i.e. he sets p

such that $q_1(p) + q_2(p) = 1$. One equilibrium involves truthful revelation: $q_i(p) = \infty$ if $p \leq 1$ and $q_i(p) = 0$ if $p > 1$. However, as Wilson (1979) has shown, there are also equilibria involving demand reduction. For example, if player 1 demands $q_1(p) = 1 - p$, then player 2 faces the residual demand curve $r_2(p) = 1 - q_1(p) = p$ and, if the price is p , his net utility is $(1 - p)p$. Consequently, if player 1 demands $q_1(p)$, then player 2 prefers the final price to be $\frac{1}{2}$. Now, if player 2 demands $q_2(p) = 1 - p$, then the final price will indeed be $\frac{1}{2}$, hence $q_2(p) = 1 - p$ is a best response against $q_1(p) = 1 - p$. By symmetry, both buyers demanding $q_i(p) = 1 - p$, therefore, is an equilibrium. This equilibrium, however, does not result in the competitive outcome: the seller's profit is only $\frac{1}{2}$.

Wilson (1979) concludes from the above example, and from more general ones in his paper, that share auctions may be less competitive than unit auctions. This insight is quite relevant for actual auctions, and it is wise to be on the outlook for possibilities of demand reduction. To give an even simpler example, assume there are 2 indivisible units for sale, and assume that there are 2 bidders that each assign a value 1 to each of these units. With demand exceeding supply, the competitive price is 1. The ascending, English auction, however has an equilibrium with revenue 0 for the seller. In this equilibrium, each buyer reduces his demand irrevocably to 1 unit at $p = 0$ and the auction stops immediately.

Up to now we have assumed that the underlying game is played only once, and the conclusions that have been derived depend on this assumption. Things may be completely different in a repeated setting. Imagine our simple (reduced) game being played repeatedly, at times $t = 1, 2, \dots$, with players discounting future payoffs with a common discount factor δ . If both players are, at each time point t , fully informed about both prices that have been asked up to that point in time, there exist the well-known collusive trigger strategy equilibria: at time t choose $p_i = 1$ if previous prices have always been 1 choose $p_i = 0$ otherwise. Such equilibria, however, rely essentially on the assumption that the entire history is commonly known: with a bit of noise and private monitoring (each player only knows his own personal history of prices and sales), such equilibria may break down, see Bhaskar and Van Damme (2002) and other papers in the recent JET-Symposium on games with imperfect private monitoring. These models are quite important for the question of whether, and in which contexts, collusion may be problematic. The reader is referred to the JET-symposium for further details.

What's the moral in the story above? I have tried to show that a very simple economic situation is actually very rich. It gives rise to a great many puzzles, which provide mathematical fun and joy. The basic situation also leads to many different models and the solution may depend strongly on a specific detail of the model. Furthermore, in some situations, multiple equilibria arise naturally, which raises the general question of how to select an equilibrium. More generally, and perhaps more importantly, there is the ques-

tion about how to select the model. Most models, however, contain relevant insights, and they all caution to the belief that competition among the few will lead to the competitive outcome. In any case, it seems quite clear that the search for a unified theory will be in vain, unless the theory itself is phrased in very general terms, which may not be very practical.

3 Differentiation, Timing and Cooperation

One message that comes from the models discussed in the previous section is that price competition indeed may be a very competitive situation, resulting in low profits for the firms involved. In such a situation, a firm might argue "if the game is so competitive, why would I play it?" To put it differently, would a firm have possibilities to change the nature of the game? In this section we discuss three possibilities for making the game less competitive: product differentiation, price leadership and cooperation. Our aim is to illustrate that such moves to change the nature of the game can also be analyzed by non-cooperative tools. Rather than focussing on the price game itself, in this case, one studies an enlarged (usually 2-stage) game in which the pure price game might arise as a subgame, but in which, because that game is so competitive, players decide to avoid it.

The product differentiation variant of this idea is due to Hotelling (1929). In his model, the set of consumers is modelled by the continuum $[0,1]$, with the consumer located at x having preferences to buy a product of characteristics x . More generally, if that consumer buys a product with characteristics y for a price p , his utility is $u(y, x) - p$. In the simplest version of the model, there are two firms who play a 2-stage game. In the first stage, they simultaneously choose product characteristics, x_1 and x_2 ; and in the second stage, given these characteristics, they simultaneously choose prices. The simple game from the previous section arises as a special case when $x_1 = x_2$. In equilibrium, the firms have to trade-off the market demand effect (*ceteris paribus*, it is most attractive to be at the center of demand, $x = \frac{1}{2}$) against the competition effect (competitive pressure is lowest if the products are maximally differentiated, i.e. if $|x_1 - x_2|$ is large) and the exact nature of the solution depends on the utility function that is specified. For the linear case analyzed by Hotelling, there is no pure strategy equilibrium; for a quadratic version, the equilibrium involves maximal differentiation; for yet other specifications, firms may decide to locate at the middle. We refer to Goeree et al. (1997) for further details and other nice results. We also note that Bester et al. (1996) have shown that multiple equilibria and coordination failures may arise, if it is not imposed exogenously that firm 1 has to locate to the left of firm 2, i.e. $x_1 \leq x_2$.

An alternative way in which too intense price competition might be avoided is if players succeed in avoiding simultaneous price setting. With firms setting prices sequentially, prices will typically be higher. A good example of this in practise is price leadership, where the market leader sets a price that is then

followed by the other players in the industry. Such price leadership exists, for example, in the Dutch gasoline market, where it is typically seen as one of the causes for the relatively high prices. More generally, the issue is one of timing of the moves, and the question is which firm will move when: who will be the leader? In joint work with Sjaak Hurkens (1996, 1999) I address this question. In these papers we ask the question which of the players will act as a leader and which one will follow, given that players have the possibility to move simultaneously or sequentially. Note that in a Cournot setting (quantity competition) it is well-known that it pays to commit: the leadership position is most preferred and the follower position least preferred, with simultaneous decisions taking an intermediate position. In price competition (with differentiated products), the situation is somewhat different: simultaneous price setting is least preferred; both players prefer there being a leader, but the follower profits most, hence, it is most advantageous to have that position.

In these papers, we introduce an asymmetry in cost functions to make the question "who will lead?" into an interesting one, and we show that the least cost firm will always emerge as the leader. Let us illustrate the model and the main idea in a simple asymmetric version of the 2-firm Bertrand game. Assume that the product is perfectly divisible, that the consumer has a marginal utility of 1 up to 1 unit and 0 thereafter, and that firm 2 has a capacity constraint of k . If firm 2 acts as a leader, firm 1 will undercut and firm 2's profit is zero. Firm 2's profit is also zero if price setting is simultaneous and in this case firm 1's profit is zero as well. If firm 1 commits to be leader, he will be undercut by firm 2, but given that firm 2 has a capacity constraint, firm 1 is not hurt that much by it. Firm 1 will simply commit to the monopoly price and profits will be $1 - k$ for firm 1 and k for firm 2. Hence, only in the case where firm 1 takes up the leadership position will profits be positive for each firm. It would be tempting to conclude from this that firms will coordinate on firm 1 acting as leader, however, such temptation should be resisted, players care only about their own payoffs and joint profits are not directly relevant. The papers with Hurkens derive the ordering of the moves from purely individualistic considerations. For the case of price competition, we consider the following 2-stage game. Players can choose a price in either period 1 or period 2, with player i being informed about the price p_j of his opponent if i chooses to move at $t = 2$ and j moves at $t = 1$. In the very special case considered here, this game can be solved by iterative elimination of dominated strategies and, not surprisingly, the outcome is that firm 1 commits to $p_1 = 1$ at $t = 1$. As said, Van Damme and Hurkens (1996) obtain the result that the efficient firm will lead in much more general settings. Still, some of the intuition is conveyed in this very simple example.

Above, we already briefly referred to the work of Edgeworth on Bertrand competition with capacity constraints. Edgeworth was also the one who introduced the Core as the concept that models unbridled competition. Shubik (1959) rediscovered this concept in the context of cooperative games and

the close relation between the Core of the cooperative exchange game and the competitive outcome was soon discovered. Weddepohl discusses the Core equivalence theorem in his first publication, Weddepohl (1971). Cooperative games are games in which the emphasis is on coalitions and in which the analyst tries to predict the outcome (i.e. which coalition might arise and how it would divide the gains from trade) without specifying the detailed rules governing the strategic interaction between the players. The motivation to study cooperative games is threefold. First of all, cooperation really is an issue; secondly one might hope that it would be possible to predict the outcome on the basis of the little information that is contained in the cooperative game, and, thirdly, as stressed by Von Neumann and Morgenstern (See Von Neumann and Morgenstern, 1953, especially Chapter V), rational players may always be expected to attempt to violate the rules of any game when it is in their interest to do so. In our most simple price competition game (which in cooperative terms amounts to the 3-person characteristic function game given by $v(S) = 1$ if $3 \in S$ and $|S| \geq 2$; and $v(S) = 0$ otherwise, where player 3 is the consumer), the core consists of a single allocation $(0,0,1)$, corresponding to the consumer buying from either producer for a price of 0, hence, the Core coincides with the competitive outcome, illustrating the Core equivalence theorem.

When, in the mid 1970s, I took my first courses in game theory with Stef Tijs, I considered the solution prescribed by the Core in the above game to be very natural. As a consequence, I was bothered very much by the fact that the Shapley value of this game was not an element of the Core and that it predicted a positive expected utility for each of the sellers. (As is well-known, the Shapley value of this game is $(1,1,4)/6$). This example and the multitude of cooperative solution concepts in general, without it being clear which one of these would apply in a specific context, induced me to turn away from Shapley value, and from cooperative theory in general, for quite a long time. It was only after I had realised that the multiplicity of solution concepts was a natural implication of trying to solve the problem at a too great level of abstraction, and that non-cooperative theory would allow one to get a better understanding of the various cooperative solution concepts, that I returned to this class of games. Of course, this position had been argued by John Nash long ago (See Nash (1953)).

Indeed, in the price competition game, coalitions arise quite naturally and, as a consequence, the Core actually provides a misleading picture. Obviously, once the sellers realize that their profits will be competed away if they do not form a cartel, they will try to form one. If they succeed in forming the cartel, they transform the situation in a bilateral monopoly in which case the negotiated price will be $\frac{1}{2}$. By symmetry, each of the sellers will get $\frac{1}{4}$ in this case. But, anticipating this, the consumer will want to form a coalition with any of the sellers, if only to prevent these sellers from entering into a cartel agreement. As Von Neumann and Morgenstern (1953) already realized, the

game is really one in which players will rush to form a coalition. (See also Schelling, 1960). But then the outcome will be determined by the coalition formation process, hence, non-cooperative modelling should focus on that process.

Let us assume that the players bump into each other at random and that, if negotiations between two players are not successful (which will not happen in equilibrium), the match is dissolved and the process starts afresh. The remaining question is what price, p , the consumer will pay to the seller if a buyer-seller coalition is formed. (By symmetry, this price does not depend on which seller the buyer is matched with.) The outcome is determined by the players' outside options, i.e. by what players can expect if the negotiations break down. The next table provides the utilities players can expect depending on the first coalition that is formed

First Coalition	Utility		
	S_1	S_2	B
S_1B	p	0	$1-p$
S_2B	0	p	$1-p$
S_1S_2	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{2}$

For the coalition S_1B , the outside option of S_1 is $\frac{1}{3}(p + \frac{1}{4})$, while the buyer's outside option is $\frac{2}{3}(1-p) + \frac{1}{6}$. (This follows since all three 2-person coalitions are equally likely to form in the next round.) The coalition loses $\frac{1}{3}(p + \frac{1}{4})$ if it does not come to an agreement, hence, it will split this surplus evenly. It follows that the price p must satisfy

$$p = \frac{1}{3}(p + \frac{1}{4}) + \frac{1}{6}(p + \frac{1}{4}),$$

Hence, $p = \frac{1}{4}$. Since all coalitions are equally likely, the expected payoff of a seller equals $\frac{1}{6}$, while the buyer's expected payoff equals $\frac{2}{3}$. The conclusion is that expected payoffs are equal to the Shapley value of the game. Furthermore, the outcome, naturally, lies outside of the Core. Of course, the exact price will depend on the details of the matching process and different processes may give rise to different prices, hence, different cooperative solution concepts. In addition, some details in the above derivation have to be filled in. All such details can be found in the literature, we refer the interested reader to Montero (2000).

What's the lesson to be drawn from the material in this Section? I would say it is that the material shows that the simple model is even richer than Section 2 indicated it to be. While in Section 2, we focused on variations in the economic data, keeping the model of price competition unchanged, we concluded in this section that, when the game itself is very competitive, players will try to find ways around it, to change the game and make the

competition less intense. Possibilities for doing so include innovation, for example, through product differentiation; strategic manipulation, for example, by influencing the timing of the game; or outright violation of the rules, by forming coalitions outside the formal rules. Clearly, one should be aware of all such possibilities, also in the policy domain. An example is provided by the Dutch UMTS-auction that took place in 2000. In that auction five licenses were for sale in a situation with five incumbents. One entrant, Versatel, participated in that auction and, just as in the above example, it had power over the incumbents: Versatel could drive up the price by staying in the auction, while it would do the incumbents a favour by stepping out; coalitions, hence, were quite natural, see Van Damme (2002).

4 Behavior

The discussion in the previous sections has been based on the assumption of rational behavior: firms and consumers maximize utility and they have no difficulty in doing so. From the experimental literature, we know that this assumption may work in some cases, but not in all: individuals may not maximize (for whatever the reason may be); or if they maximize, they need not maximize the utility function that we imputed; all kinds of other arguments might play a role. This raises the question of how well the rationality assumption works in the special, simple Bertrand game that we are considering in this paper.

Surprisingly, even though we know something about the answer to this question, we know remarkably little (See Holt (1995) for an overview.). In any case, there is some good news: the concerns for fairness that play such an important role in ultimatum bargaining experiments seem to play much less of a role in competitive price setting situations. On the other hand, there is evidence that experimental subjects realise that it is a sensible strategy to try to get out of the competitive situation. The evidence comes, among others from experiments conducted at Tilburg University, which are reported in Dufwenberg and Gneezy (1998). In what follows, we briefly describe their experiment.

The game they analyse, is based on the simple example with which we started this paper, with only the two sellers active. They bid integers from the set $\{2, 3, \dots, 100\}$ and the one who chooses the lowest number, say n , gets ng where g is a parameter, a certain number of guilders, with ties being split. Clearly, in the unique Nash equilibrium of this game, each player chooses 2 and each seller receives g . In the experimental implementation of this game, there were 6 groups of 12 students and each group played 10 rounds of the game. In each round, 6 pairs of players were randomly formed, with no student meeting another student more than once. Three treatments were distinguished: in the high information treatment, H , before round $t + 1$ each player was informed about all bids made in round t ; in the medium informa-

tion treatment, *M*, each player got to see all the winning bids, and in the low information treatment, *L*, each player was only informed about his own personal pay-off at the end of the round. Hence, the treatments correspond to various degrees of transparency about the market. For fully rational players, the amount of information should, however, be irrelevant as the equilibrium of the game does not depend on it.

How to play in such a game? What did the outcomes look like? How did the outcome depend on the information feedback? It turns out that in all sessions, behavior in round 1 is very much different from the Bertrand outcome. This actually should not be too surprising: players have to learn first to understand the game. Hence, it is more instructive to look at the data from round 10, but also then the outcomes may be different. The authors conclude: "it seems that in the treatments *M* and *L* the winning bids converged (almost) to the Nash equilibrium, whereas in treatment *H* there was no tendency at all of convergence to the Nash equilibrium". Formally, using the Mann-Whitney *U* test, the authors reject (at the 5% level) that the prices in the *H*-treatment are drawn from the same distribution as was observed in the *M*- or *L*-treatment. Clearly, it is encouraging that, in the *L*- and *M*-treatments, prices converge quickly to the Nash equilibrium, but what is going on in the *H*-treatment? The authors argue that, when losing bids are observed, such bids may be used to signal desirable outcomes. Indeed, some of the behavior in their data is consistent with such signalling behavior. Nevertheless, the data appear somewhat puzzling: even if signalling might be a good thing during the game, it seems that there is little to gain from it in the last round. So why don't prizes suddenly drop when players know it is the last round?

Despite the above question, there is the important observation by Dufwenberg and Gneezy that a more transparent market is not necessarily more competitive. On the contrary, more transparency may make signalling easier and tacit collusion more likely. A similar message comes from other recent experimental research, such as Huck et al (1999) and Offerman et al (2001). Of course, a similar message was already conveyed by the older experimental work, such as that reported in Fouraker and Siegel (1953). Nevertheless, this message stands in stark contrast to the claim that the Internet will make markets both more transparent and more competitive (Van Damme and Delaert (2001)) and it seems that more systematic research in this area would be worthwhile.

The main message from this section is simple. In some experiments with the Bertrand model, behavior converges quickly to the Nash equilibrium. In others, it does not. Some variables, which according to standard theory should not matter, do matter in practice. More work is to be done.

5 Conclusion

Let me try to summarize the argument that I have been trying to make. It is no doubt important to try to understand price competition among the few; will the outcome be close to the competitive outcome, or not? Some 115 years ago, Bertrand proposed a very simple model of price competition. Even though the model is simple, it is also very rich in the sense that it allows elaborations in many directions. It is also intellectually stimulating, and great fun to analyse and try to understand the variations on Bertrand's basic model. The richness of the structure has advantages and disadvantages. The great advantage is that, given an actual market, a puzzling observation or a paradox, it will be reasonably easy to find a Bertrand-like model that captures some essential element of that situation, and that is at the same time reasonably well understood, so that the model may be expected to illuminate the market or puzzle. The drawback, of course, is that it will be difficult to have a unified theory, unless one is satisfied with a very general, abstract theory, or with a theory that is not neatly formalised, that is, a theory that is "just" a collection of connected insights.

Personally, I am happy with the Bertrand theory being loose, abstract and not completely formalised; in my view, the advantages mentioned above outweigh the disadvantages several times. In this sense indeed relevance may be preferred to rigour, at least if rigour is interpreted as insistence on a clear demarcation of in which contexts a specific insight will apply. Such a sharp boundary may not be possible, and, instead, one may be satisfied with relying on judgment in this respect. Of course, at the same time, one should insist on full rigour within a certain model. Reviewing the above material, I believe there are several insights there, such as the competitiveness of naked price competition, the desire of players to avoid too intense competition, hence, the incentives to change the nature of the game, the multiplicity of supply function equilibria (an insight especially relevant for the electricity market), the importance of market transparency, and the relation, up to now imperfectly understood, between transparency and competition (a question extremely relevant for Internet markets), and the insight that whether or not a Nash equilibrium is attained may depend on variables that are considered irrelevant in the standard theory.

The above point of view implies that a good economist probably should both have good powers of observation (in order to be able to detect puzzles), good judgement (so as to extract essential elements and build relevant models), and good mathematical skills (in order to structure, formalise and solve these models). In this respect, the strict separation that exists in the Netherlands between economics and econometrics is to be regretted, for it is unlikely to produce all-round economists that master all those skills. Obviously, what is unlikely is not impossible, and Weddepohl might be one of these exceptions that prove the rule. Another piece of good news is that also in the Netherlands, the sharp separation between the two fields is van-

ishing, with some of Weddepohl's students being active on both sides and, hence, actively contributing to tearing down the wall. The reader will probably be able to judge this himself, based on their contributions in this volume. As for Weddepohl's judgement, let me simply refer the reader back to the quote from the introduction, and let me note that indeed, because of Dutch (policy) economists being unfamiliar with oligopoly models, there has been a tendency to view oligopolistic markets as being competitive, with all the resulting disappointment.

References

1. Baye, Michael and John Morgan (1999), A folk theorem for one-shot Bertrand Games, *Economics Letters* 65, 59-65
2. Bester, Helmut, André de Palma, Wolfgang Leininger, Elu von Thadden and Jonathan Thomas (1996), A non-cooperative analysis of Hotelling's Location Game, *Games and Economic Behavior* 12, 165-186.
3. Bhaskar, V. and Eric van Damme (2002), Moral Hazard and Private Monitoring, *Journal of Economic Theory*, Forthcoming.
4. Brandenburger, Adam and Jerome Keisler (2000), Epistemic conditions for iterated admissibility, *Discussion Paper*, Harvard Business School.
5. Canoy, Marcel and Claus Weddepohl (1995) Alternative conjectures in a Bertrand-Edgeworth Model, *European Journal of Political Economy* 11, 577-598.
6. Damme, Eric van (2001a) Marktwerking Vereist Maatwerk *Maandschrift Economie* 65(3) 185-208.
7. Damme, Eric van (2001b) MDW tegen het Licht, *Regelmaat* 2001/3, 87-97.
8. Damme, Eric van (2002) The European UMTS-Auctions, *European Economic Review*, Forthcoming.
9. Damme, Eric van and Sjaak Hurkens (1996), Endogenous Price Leadership, *Discussion Paper* nr. 96115, CentER, Tilburg University.
10. Damme, Eric van and Sjaak Hurkens (1999), Endogenous Stackelberg Leadership, *Games and Economic Behavior* 28, 105-129.
11. Damme, Eric van and Benedict Dellaert (2001) e-Economie: ICT en Marktwerking, in Rick van der Ploeg and Chris Veenemans (eds.) *De invloed van ICT op maatschappij en overheid*, Infodrome, 203-285.
12. Dasgupta, Partha and Eric Maskin (1986) The existence of equilibria in discontinuous games: II, Applications, *Review of Economic Studies* 53, 27-41.
13. Dastidar, K. (1995) Comparing Cournot and Bertrand in a Homogeneous product market, *Journal of Economic Theory* 75, 205-212.
14. Diamond (1971) A model of price adjustment, *Journal of Economic Theory* 3, 156-169.
15. Dufwenberg, Martin and Uri Gneezy (1998), Price Competition and Market Concentration: An Experimental Study, *Discussion Paper* nr. 9827, CentER, Tilburg University.
16. Dufwenberg, Martin and Mark Stegeman (2001), When order matters for iterated strict dominance, *Econometrica*, Forthcoming

17. Edgeworth, Francis Ysidio (1897) The Pure Theory of Monopoly, reprinted in William Baumol and Stephan Goldfeld *Precursors in Mathematical Economics: An Anthology*, London School of Economics, 1968.
18. Fouraker, D. and Sidney Siegel (1953) *Bargaining behaviour*, McGraw Hill
19. Goerree, Jacob, Simon Anderson and Roald Ramer (1997) Location, Location, Location, *Journal of Economic Theory* **77**, 102-127.
20. Holt, Charlie (1995) Industrial Organization: A survey of Laboratory Research, Chapter 5 in John Kagel and Alvin Roth (eds.), *Handbook of Experimental Economics*, Princeton University Press, Princeton, NJ 1995.
21. Hotelling, Harald (1929) Stability in Competition, *Economic Journal* **39**, 41-57.
22. Huck, Stephan, Hans-Theo Normann and Jörg Oechsler (1999) Learning in Cournot Oligopoly: An Experiment, *Economic Journal* **109**, C80-C95
23. Marschak, Thomas and Reinhard Selten (1974), Equilibrium with Price Making Firms, *Lecture Notes in Economics and Mathematical Systems* **91**, Springer Verlag.
24. Montero, Maria (2000) *Endogeneous Coalition Formation and Bargaining*, PhD thesis, CentER, Tilburg University.
25. Nash, John (1953) Two-person Co-operative Games, *Econometrica* **21**, 128-140.
26. Neumann, John von and Oskar Morgenstern (1953) *Theory of Games and Economic Behaviour*, Princeton University Press, Princeton, NJ (First edition, 1943).
27. Offerman, Theo, Jan Potters and Joep Sonnemans, Imitation and belief learning in an oligopoly experiment, Discussion paper, CREED University of Amsterdam and CentER Tilburg University. Forthcoming in *Review of Economic Studies*.
28. Schelling, Thomas (1960) *The Strategy of Conflict*, Harvard University Press, Cambridge, MA.
29. Shubik, Martin (1959) Edgeworth market games, in A.W. Tucker and R.D. Luce (eds.) *Contributions to the theory of games IV; Annals of Mathematics Studies* **40**, Princeton University Press, Princeton, NJ.
30. Simon, Leo and Max Stinchcombe (1995) Equilibrium Refinement for Infinite Normal-Form Games, *Econometrica* **63**, 1421-1443.
31. Weddepohl, Claus (1971) Concurrentie-Evenwicht en de Core van een Directe Markt, *Maandschrift Economie* **35**, 427-448.
32. Weddepohl, Claus (1975). Concurrentie in Geval van Afnemende Gemiddelde Kosten, *Maandschrift Economie* **40**, 87-102.
33. Weddepohl, Claus (1991) Op Zoek naar Evenwicht in een Bertrand Model, in Bouma, J. e.d. (eds.), *Samenspel, Opstellen aangeboden aan Prof.dr. G.F.W. Pikkemaat*, Wolters Noordhof, Groningen, 85-102.
34. Wilson, Robert (1979) Auctions of Shares, *Quarterly Journal of Economics* **93**, 675-689