

Oligopoly Games for Use in the Classroom and Laboratory

Ate Nieuwenhuis^{1,*}

Eindhovenpad 3B, 3045 BK Rotterdam, The Netherlands

Abstract

To illustrate that a Nash Equilibrium results from a flawed attempt to solve a game, this article studies two extensions of the classic oligopoly model of Cournot. The common cost function is quadratic, and the (still linear) inverse demand functions allow of differentiated goods. The industry has a maximal-profit set that is characterised by a constant profit-output ratio, independent of the number of firms and the slopes of the marginal-cost function and inverse demand functions. The maximal-profit set and the Pareto optimal set—which is the model’s solution—have a number of points in common, but in general do not coincide. The choice of parameters is discussed and six model variants are analysed numerically; in five of them, the incentives for merger according to noncooperative game theory are at odds with the rationale of economics. Some comments are made on the use of the model in experimental economics.

Keywords: Bertrand, Cournot, Collusion, Oligopoly, Game theory, Vector maximisation

JEL: C61, C70, D21, D43, L11, L13

No one has the right, and few the ability, to lure economists into reading another article on oligopoly theory without some advance indication of its alleged contribution. The present paper accepts the hypothesis that oligopolists wish to collude to maximize joint profits. It seeks to reconcile this wish with facts, such as that collusion is impossible for many firms and collusion is much more effective in some circumstances than in others.

G. J. Stigler (1964, p. 44)

1. Introduction

There is something weird about the notion of Nash Equilibrium. Although game theory has been designed to deal with decision problems in which the effects of each player’s actions on the outcomes for the other players cannot be ignored, noncooperative game theory does just that: the “first-order conditions” for a Nash Equilibrium ignore the partial cross-derivatives of the payoff functions. Ignoring the cross-derivatives is in fact a grave mathematical error: it amounts to treating the actions inconsistently and leads to contradictions. The two-page note Nieuwenhuis (Unpublished results, 2018) merely underlines that the conditions for a Nash Equilibrium are incompatible with the first-order approximations of the payoff functions. Nieuwenhuis (Unpublished results, 2017) exposes the error in noncooperative game theory at greater length, arguing that a game is a vector maximisation problem from which a Nash Equilibrium derives by treating variables erroneously as constants in certain places, and illustrates the fallacy of Nash Equilibrium with three examples, in oligopoly theory, in the Prisoner’s Dilemma, and in dynamic general equilibrium models with imperfect competition and rational expectations. As an example within the first example, the paper analyses the classic duopoly model of Cournot (1927)—homogeneous product, linear inverse demand function and cost function—extended to the case of a common quadratic cost function.

*E-mail address: ate_1949@hotmail.com. Phone: +31 (0)6 37 34 39 39.

¹No affiliation. This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors. Declaration of interests: none.

The present article analyses a further extension of the “classic case” to illustrate that a Nash Equilibrium results from a mathematically flawed attempt to solve a game, thus insisting on the removal of noncooperative game theory from the textbooks on microeconomic theory. It considers oligopolies with any number of firms that produce varieties of some good (one homogeneous good included as a special case), sharing a quadratic cost function. The inverse demand functions are still linear; for simplicity’s sake, the specification is symmetric in the way it treats the goods. Like Stigler’s, the article accepts the hypothesis that firms wish to collude to maximise profits, but it takes a different turn. It stresses that collusion *is* in fact the profit maximising strategy of the firms in the artificial world of the model, as this world does not contain any impediments for collusion. It defines the *maximal-profit set* of an industry as the set of points where industry profit is maximal along rays through the origin, and shows that it is characterised by a constant profit-output ratio, independent of the number of firms and of the slopes of the marginal-cost function and inverse demand functions. The maximal-profit set and the Pareto Optimum—which is the model’s solution—have a number of points in common, but in general do not coincide. The article discusses how parameter values may be chosen so as to obtain model variants with attractive features for use in the laboratory, and presents numerical examples to highlight the differences between and similarities of the outcomes of six model variants; in five variants, the incentives for merger according to noncooperative game theory appear at odds with the tenet of rational, optimising behaviour. Finally, it comments on the use of this class of models in the economics laboratory.

2. The model

Consider an industry with profit maximising firms $i, i = 1, \dots, I$, each of which produces a variety of some good. The firms have identical quadratic cost functions and linear marginal-cost functions,

$$c(q_i) = c_1 q_i + c_2 q_i^2, \quad (1)$$

$$mc(q_i) = c_1 + 2c_2 q_i, \quad (2)$$

with q_i the quantity of the variety produced by Firm i , and where $c_2 \gtrless 0$, so that marginal cost may be increasing, constant or decreasing. The last case requires $c_1 \gg 0$ for marginal cost to be positive over a certain range of output.

The inverse demand functions, too, are linear:

$$p_i = f_i(q_1, \dots, q_I) = d_0 - dQ - (t - d)q_i, \quad i = 1, \dots, I, \quad (3)$$

where $Q := \sum_j q_j$. The entity Q will serve as an indicator of industry output. The parameters satisfy $d_0 > c_1, 0 < d \leq t$. The condition $d_0 > c_1$ ensures that the prices of the goods exceed marginal cost when the firms produce nothing at all. When $t > d$, the goods are imperfect substitutes. The symmetric specification of the inverse demand functions, next to the common cost function, implies that most outcomes of interest are points on the ray $q_1 = q_2 = \dots = q_I$ (the axis of symmetry of the model), which keeps the analysis simple.

With differentiated goods, that is when $t > d$, the ordinary demand functions are

$$q_i = g_i(p_1, \dots, p_I) = \delta_{0,I} + \delta_I P - (\tau_I + \delta_I) p_i, \quad i = 1, \dots, I, \quad (4)$$

$$\delta_{0,I} = \frac{d_0}{(I-1)d+t}, \quad I \geq 1,$$

$$\tau_I = \frac{(I-2)d+t}{((I-1)d+t)(t-d)}, \quad I \geq 1,$$

$$\delta_I = \frac{d}{((I-1)d+t)(t-d)}, \quad I \geq 2,$$

where $P := \sum_j p_j$. The parameters satisfy $\delta_{0,I} > 0, 0 < \delta_I < \tau_I$.

The systems of inverse or ordinary demand functions may also be *partially* inverted to a system of *mixed* demand functions, with the prices of some goods and the quantities of the other goods as left-hand side variables. In general, there are 2^I equivalent systems of demand functions. The present, symmetric specification reduces the number of analytically distinct cases to $I + 1$, one case for each number of firms that use the prices of their products as instruments.

The profit function of Firm i is revenue minus cost, $w_i(p_i, q_i) := p_i \cdot q_i - c(q_i)$. The profit functions of all firms must be maximised subject to the constraint of the cost function (1) and to the constraints of the inverse demand functions (3) or, equivalently, the ordinary demand functions (4) or some system of mixed demand functions. These are *extensive* forms of the problem, all equivalent to one another. The I problems of profit maximisation are interdependent, because the arguments of the maximands are interrelated through the demand functions: they constitute a vector maximisation problem² or game. The game, in any extensive form, may be treated with the methods of De Finetti (1937a,b)³ or Kuhn and Tucker (1950, Section 6).⁴ The alternative and more common approach is to use a system of demand functions for eliminating one half of the number of arguments from the maximands to arrive at a *normal* form of the game; in each normal form, the arguments of the maximands may be varied independently of one another. Every normal form, combined with the matching system of demand functions, is equivalent to an extensive form; hence, all normal forms are equivalent to one another. The choice of which variables to eliminate is arbitrary in the sense that, *with a correct mathematical treatment*, it does not affect the solutions for the prices and quantities.⁵ A game theorist who attacks a normal form must secure that her proposed solution respects the equivalence of the normal forms.

The choice of which variables to eliminate is not arbitrary in every aspect. In fact, there is a most convenient choice, which is to eliminate the prices. It yields easy to understand expressions, because the dimensions of the parameters of the cost function and the inverse demand functions agree: c_1 has the same dimension as d_0 does, c_2 has the same dimension as d and t do. The profit functions *in quantity space* are

$$v_i = u_i(q_1, \dots, q_I) = (d_0 - c_1 - dQ)q_i - (t - d + c_2)q_i^2, \quad i = 1, \dots, I. \quad (5)$$

They differ from the profit functions of the classic case only if $t - d + c_2 \neq 0$. However, even if $t - d + c_2 = 0$, the ordinary demand functions and, hence, the profit functions *in price space* do differ when $t > d$. More generally, every triplet (t, d, c_2) that satisfies $t - d + c_2 = z$ yields the same profit functions in quantity space but not in price space. This property of the model results from the linearity of the marginal-cost function and the inverse demand functions. It means that to every increase of the slope of the marginal-cost function there corresponds a decrease of the slopes of the perceived marginal-revenue functions (for which see below) so as to yield the same profit functions in quantity space. The outcomes of model variants in Section 4 will illustrate the property.

The profit of Firm i is zero if $q_i = 0$ or else if $p_i = c(q_i)/q_i$. The latter condition defines a (hyper)plane the nonnegative segment of which will be called the zero-profit plane of the firm. Let Z be the intersection of the ray $q_1 = q_2 = \dots = q_I$ with the zero-profit plane of any one of the firms. Then Z is either the centroid of the common zero-profit plane (if $t - d + c_2 = 0$) or else the intersection of the zero-profit planes, the unique zero-profit point. All its coordinates equal

$$qz = \frac{d}{(I-1)d + t + c_2}k, \quad k := \frac{d_0 - c_1}{d}.$$

Other quantities of interest, too, will be expressed as fractions of k .

The solution of a game is the Pareto Optimum. The Pareto Optimum of the oligopoly game is given in Appendix B. It is a (hyper)surface in the nonnegative orthant with the ray $q_1 = q_2 = \dots = q_I$ as its axis of symmetry. In the main text the focus is on its centroid *Col* (for Collusion Point), which is compared to two “equilibria” that are also points on the axis of symmetry. Appendix B shows that in some cases *Col* is the Joint Profit Maximum, in other cases it is the point of the Pareto Optimum where the joint profit is minimal. To find the coordinates of *Col*, note that along the ray $q_1 = q_2 = \dots = q_I$ the profit of each firm is a quadratic function of the common quantity q , and that the zeros of the parabola are at the origin and at Z . Therefore *Col* is halfway between the origin and Z .

Two failed attempts to solve the vector maximisation problem constituted by the simultaneous maximisation of the profit functions are those by Cournot (1927) and by Bertrand (1883). The common trait of the attempts is that they condition on *endogenous* variables in the derivation of first-order conditions for maximal profits of the firms, and they do so in such a way that even the *possibility* of undoing the conditioning in a subsequent step of the optimisation

²Elsewhere I have used the term *simultaneous maximum problem*.

³Translations of these articles are available in Nieuwenhuis, A., 2017. Simultaneous Maximisation in Economic Theory. <https://dx.doi.org/10.13140/RG.2.2.11452.95360>.

⁴See Nieuwenhuis (Unpublished results, 2017, Section 2.3). Noncooperative game theory cannot deal with a game in extensive form.

⁵Else the eliminations would not even have been permissible.

process is lacking. In a Cournot Oligopoly, each firm chooses its quantity while conditioning on the quantities of its rivals; in a Bertrand Oligopoly, each firm chooses its price while conditioning on the prices of its rivals. The firms in a Bertrand Oligopoly perceive the marginal-revenue functions

$$mr_i^B = p_i - \frac{1}{\tau_I} g_i(p_1, \dots, p_I) \quad (6a)$$

$$= f_i(q_1, \dots, q_I) - \frac{1}{\tau_I} q_i. \quad (6b)$$

When $t = d$ (one homogeneous product), $1/\tau_I = 0$ so that $mr^B = p$. The firms in a Cournot Oligopoly perceive the marginal-revenue functions

$$mr_i^C = f_i(q_1, \dots, q_I) - tq_i \quad (7a)$$

$$= p_i - tg_i(p_1, \dots, p_I). \quad (7b)$$

Subsequently equating marginal revenue according to the functions (6b) or (7a) to marginal cost (2) yields the “reaction functions” in implicit form. The Bertrand Equilibrium B and the Cournot Equilibrium C are at the intersection of the ray $q_1 = q_2 = \dots = q_I$ with any one of the appropriate “reaction functions.”

The concept of “reaction function” plays no role in the derivation of the Pareto optimal set. Still, the Collusion Point may be derived as the Bertrand Equilibrium and Cournot Equilibrium are. Due to the model’s symmetry (see Appendix A.1), the thought experiment of varying the quantities in unison or, alternatively, of varying the prices in unison yields the firms at Col as if they perceive the marginal-revenue functions

$$mr_i^{Col} = f_i(q_1, \dots, q_I) - ((I-1)d+t)q_i \quad (8a)$$

$$= p_i + \frac{1}{(I-1)\delta_I - \tau_I} g_i(p_1, \dots, p_I). \quad (8b)$$

The expressions (8a) and (8b) are equivalent, meaning that the Collusion Point is unaffected by the choice of instruments. In fact, the Pareto optimal set is unaffected by nonsingular transformations of variables (as is the solution of every optimisation problem). The Cournot Equilibrium and Bertrand Equilibrium, by contrast, two examples of a Nash Equilibrium, will differ from one another. In general, the Nash Equilibria of all normal forms will be different. The dependence of a Nash Equilibrium on the *arbitrary* choice of instruments points to the flawed treatment of the vector maximisation problem, and is a major defect of noncooperative game theory as a theory of rational behaviour. Appendix A.2 explains how a Nash Equilibrium results from unwarranted conditioning on endogenous variables.

Table 1 gives a number of outcomes. The prices are always a weighted average of c_1 and d_0 . The outcomes of the classic case are the limits of the outcomes of the general case when the number of firms grows without bound: the effects of nonconstant marginal cost and of product differentiation become negligible. A noteworthy feature of the Collusion Point is the constant profit-output ratio,⁶

$$\frac{v_{Col}}{q_{Col}} = \frac{d_0 - c_1}{2},$$

independent of I , the number of firms, and of c_2 , d and t , the slopes of the marginal-cost function and inverse demand functions. In fact, this outcome is a special case of a more general result. Define the *maximal-profit set* of an industry as the set of points where industry profit is maximal along the rays $\alpha_1 q_1 = \alpha_2 q_2 = \dots = \alpha_I q_I$, $\alpha_i \geq 0$ for all i . There holds

FACT: *The maximal-profit set of every oligopoly in this paper is the set of points where the profit-output ratio of the industry equals $(d_0 - c_1)/2$.*

⁶For the quadratic form $ax^2 + bx$, the ratio of the extremum to its argument is $b/2$.

Table 1: Main outcomes

	General formula	The classic case
Zero-profit point		
– quantity $\div k$	$\frac{d}{(I-1)d+t+c_2}$	$\frac{1}{I}$
– price	$\frac{((I-1)d+t)c_1+c_2d_0}{(I-1)d+t+c_2}$	c_1
Bertrand Equilibrium		
– quantity $\div k$	$\frac{d}{(I-1)d+t+\tau_I^{-1}+2c_2}$	$\frac{1}{I}$
– price	$\frac{((I-1)d+t)c_1+(\tau_I^{-1}+2c_2)d_0}{(I-1)d+t+\tau_I^{-1}+2c_2}$	c_1
– profit	$(\tau_I^{-1}+c_2)q_B^2$	0
Cournot Equilibrium		
– quantity $\div k$	$\frac{d}{(I-1)d+2(t+c_2)}$	$\frac{1}{I+1}$
– price	$\frac{((I-1)d+t)c_1+(t+2c_2)d_0}{(I-1)d+2(t+c_2)}$	$\frac{Ic_1+d_0}{I+1}$
– profit	$(t+c_2)q_C^2$	dq_C^2
Collusion point		
– quantity $\div k$	$\frac{d}{2((I-1)d+t+c_2)}$	$\frac{1}{2I}$
– price	$\frac{((I-1)d+t)c_1+((I-1)d+t+2c_2)d_0}{2((I-1)d+t+c_2)}$	$\frac{c_1+d_0}{2}$
– profit	$((I-1)d+t+c_2)q_{Col}^2$	Idq_{Col}^2

The maximal-profit set and the Pareto optimal set have the Collusion Points of the included oligopolies in common, but in general do not coincide; Appendix B delves deeper into this matter. Still it is true that, within the Pareto optimal sets of the oligopolies with the same value of $d_0 - c_1$, the joint profit is largest where the joint output is largest. In Section 4 the result will be used in the comparison of the outcomes of six model variants.

3. Choice of parameters: the classic case

Numerical examples help to reveal the information that the expressions in Table 1 contain. The present section discusses the choice of values for the parameters c_1 , d_0 and d of the classic case, because it takes a central place. The next section will show the impact of nonconstant marginal cost and product differentiation on the outcomes.

Conveniently, choosing values for the parameters can be done in three steps. First, scale the quantities through the choice of k ($:= (d_0 - c_1)/d$). Second, fix the price at the Collusion Point through the choice of $c_1 + d_0$. Third, make an assumption about the elasticity of the inverse demand function at the Collusion Point to identify c_1 , d_0 and d .

Table 2 gives the unchanging pattern in the outcomes for 1–5 firms. The table must be transformed into a numerical example with certain desirable properties, which make the model suitable for application in the laboratory. One wish

Table 2: The classic case: pattern in the outcomes

Number of firms (I)	1	2	3	4	5
Quantities $\div k$					
$-q_Z$	1	1/2	1/3	1/4	1/5
$-q_B$	1/2	1/2	1/3	1/4	1/5
$-q_C$	1/2	1/3	1/4	1/5	1/6
$-q_{Col}$	1/2	1/4	1/6	1/8	1/10
Prices					
$-p_Z$	c_1	c_1	c_1	c_1	c_1
$-p_B$	$\frac{c_1 + d_0}{2}$	c_1	c_1	c_1	c_1
$-p_C$	$\frac{c_1 + d_0}{2}$	$\frac{2c_1 + d_0}{3}$	$\frac{3c_1 + d_0}{4}$	$\frac{4c_1 + d_0}{5}$	$\frac{5c_1 + d_0}{6}$
$-p_{Col}$	$\frac{c_1 + d_0}{2}$	$\frac{c_1 + d_0}{2}$	$\frac{c_1 + d_0}{2}$	$\frac{c_1 + d_0}{2}$	$\frac{c_1 + d_0}{2}$

is that the Bertrand Equilibrium B , Cournot Equilibrium C and Collusion Point Col be clearly separated; given the fixed ratios between the outcomes, this property can only be obtained by choosing the scale “sufficiently” large. A sufficiently large scale also helps to constrain the relative deviation from the true outcomes introduced by rounding these to the nearest integer, a practice that seems advisable in the laboratory. On the other hand, the figures must not be “unduly” large. It would also be nice for the outcomes to have a somewhat “realistic” flavour.

The admittedly vague desiderata leave ample room for other considerations. The choice has been further limited to parameters that yield integer outcomes for the cases of 1–5 firms. As to the quantities, it requires k to be a multiple of 120 ($= 2^3 \cdot 3 \cdot 5$). It does not seem necessary to choose k larger than 120. As to the prices, it requires $c_1 + d_0$ to be a multiple of 60 ($= 2^2 \cdot 3 \cdot 5$) if $c_1 = 0$. The value of 60 implies $p_{Col} = 30$. Given the need to choose c_1 well in excess of 0 (to allow of decreasing marginal cost), this value leaves a rather small interval for the prices. So let’s double it.

Table 3: The classic case: a numerical example

I	1	2	3	4	5
Quantities					
$-q_Z$	120	60	40	30	24
$-q_B$	60	60	40	30	24
$-q_C$	60	40	30	24	20
$-q_{Col}$	60	30	20	15	12
Prices					
$-p_Z$	30	30	30	30	30
$-p_B$	60	30	30	30	30
$-p_C$	60	50	45	42	40
$-p_{Col}$	60	60	60	60	60

Lastly, the wish for a somewhat “realistic” flavour. Observe that the elasticity of the inverse demand function at the Collusion Point, e_{Col} , is given by $(c_1 - d_0)/(c_1 + d_0)$. A value of c_1 close to d_0 yields a value of e_{Col} close to zero and hence a low markup of price over (marginal) cost, whereas a value of c_1 close to zero yields a value of e_{Col} close to -1 and hence a high markup. To steer clear of both extremes, choose $e_{\text{Col}} = -0.5$, which yields a markup of 1 and implies $d_0 = 3c_1$. The ensuing parameter values are $c_1 = 30$, $d_0 = 90$ and $d = 0.5$. Table 3 gives the outcomes for the quantities and prices with this choice of parameters. Note that the parameter d may be used to scale the quantities. For example, halving the value of d doubles all quantities: the change represents a pure demand shift, with the same relative rise of demand at every price.

4. Six variants

The stage is now set for model variants with nonconstant marginal cost ($c_2 \neq 0$) and/or differentiated goods ($t > d$). It is convenient to choose c_2 and $t - d$ in proportion to d : this practice yields all quantities as fractions of k that are independent of d , so that a change of d still represents a pure demand shift. For c_2 the values of $d/4$, 0 and $-d/4$ are considered, for $t - d$ the values of 0 and $d/4$. Table 4 gives the labels of the six model variants, or *industries*.

Table 4: Six model variants

Nature of the good(s)	Marginal cost		
	Increasing	Constant	Decreasing
	$(c_2 = d/4)$	$c_2 = 0$	$(c_2 = -d/4)$
Homogeneous: $t - d = 0$	HoIn	HoCon	HoDe
Heterogeneous ($t - d = d/4$)	HeIn	HeCon	HeDe

One element in the discussion of the outcomes is the comparison within an industry across the numbers of firms, which have been treated as exogenous so far. Such a comparison naturally leads to the question what the assumption of joint profit maximisation implies. To answer it, the Tables 5–10 contain, next to the prices, the outcomes for total output Q and total profit U of the industries. The discussion uses the results on the Pareto optimal set in Appendix B, specifically that industry profit is largest where total output is largest. To avoid the infinite and infinitesimal, it is assumed that there is some non-zero minimal firm size and that firms can enter an industry only at the minimal size.

Consider first the outcomes of the classic case, Industry HoCon, in Table 7. A Bertrand duopoly produces twice the monopoly quantity, which is a generic outcome of the classic case. The price equals marginal cost at half the monopoly price (a consequence of the specific choice of parameters), and profits are down to zero. A further increase of the number of firms changes neither the total output nor the price. As the number of firms rises, the Cournot Equilibrium moves gradually from the monopoly outcome towards the Bertrand Equilibrium and the zero-profit competitive outcome; the increase from one firm to ten firms closes most of the gap between the monopolistic and competitive outcomes.⁷ Actually, the Cournot Oligopoly owes its popularity to this gradual transition, as it agrees with the intuition of many economists. Another way of looking at the same pattern, however, is that the firms in the Cournot Oligopoly perceive an incentive to merge or to collude: industry profit rises as the number of firms declines. But the true profit maximising strategy for the firms is in fact to charge the monopoly price and to jointly produce the monopoly output. As shown in Appendix B, in this industry the Pareto optimal set is flat and coincides with the maximal-profit set. The optimal profit-output ratio being the same over the whole range of firm sizes and numbers of firms, the joint profit is constant, too: the assumption of joint profit maximisation does not select a specific size distribution of firms. Nonconstant marginal cost and/or product differentiation usually change this outcome, as will appear below.

⁷For all three models in all six industries, most outcomes for the decapoly are close to the limit values.

Table 5: Industry HoIn

I	1	2	3	4	5	10
Q_Z	96	107	111	113	114	117
Q_B	48	96	103	107	109	114
Q_C	48	69	80	87	92	104
Q_{Col}	48	53	55	56	57	59
p_Z	42	37	35	34	33	31
p_B	66	42	39	37	35	33
p_C	66	56	50	46	44	38
p_{Col}	66	63	62	62	61	61
U_B	1440	576	441	356	298	163
U_C	1440	1469	1333	1190	1065	681
U_{Col}	1440	1600	1662	1694	1714	1756

Table 6: Industry HeIn

I	1	2	3	4	5	10
Q_Z	80	96	103	107	109	114
Q_B	40	75	88	95	99	109
Q_C	40	60	72	80	86	100
Q_{Col}	40	48	51	53	55	57
p_Z	40	36	34	33	33	31
p_B	65	48	43	40	38	34
p_C	65	56	51	48	45	39
p_{Col}	65	63	62	62	61	61
U_B	1200	984	781	645	548	312
U_C	1200	1350	1296	1200	1102	750
U_{Col}	1200	1440	1543	1600	1636	1714

Table 7: Industry HoCon

I	1	2	3	4	5	10
Q_Z	120	120	120	120	120	120
Q_B	60	120	120	120	120	120
Q_C	60	80	90	96	100	109
Q_{Col}	60	60	60	60	60	60
p_Z	30	30	30	30	30	30
p_B	60	30	30	30	30	30
p_C	60	50	45	42	40	35
p_{Col}	60	60	60	60	60	60
U_B	1800	0	0	0	0	0
U_C	1800	1600	1350	1152	1000	595
U_{Col}	1800	1800	1800	1800	1800	1800

Table 8: Industry HeCon

I	1	2	3	4	5	10
Q_Z	96	107	111	113	114	117
Q_B	48	89	100	105	108	114
Q_C	48	69	80	87	92	104
Q_{Col}	48	53	55	56	57	59
p_Z	30	30	30	30	30	30
p_B	60	40	36	34	33	32
p_C	60	51	47	44	42	37
p_{Col}	60	60	60	60	60	60
U_B	1440	889	598	449	360	180
U_C	1440	1469	1333	1190	1065	681
U_{Col}	1440	1600	1662	1694	1714	1756

Table 9: Industry HoDe

I	1	2	3	4	5	10
Q_Z	160	137	131	128	126	123
Q_B	80	160	144	137	133	126
Q_C	80	96	103	107	109	114
Q_{Col}	80	69	65	64	63	62
p_Z	10	21	25	26	27	28
p_B	50	10	18	21	23	27
p_C	50	42	39	37	35	33
p_{Col}	50	56	57	58	58	59
U_B	2400	-1600	-864	-588	-444	-199
U_C	2400	1728	1322	1067	893	490
U_{Col}	2400	2057	1964	1920	1895	1846

Table 10: Industry HeDe

I	1	2	3	4	5	10
Q_Z	120	120	120	120	120	120
Q_B	60	109	116	118	119	120
Q_C	60	80	90	96	100	109
Q_{Col}	60	60	60	60	60	60
p_Z	15	23	25	26	27	29
p_B	53	29	27	27	28	29
p_C	53	45	41	39	38	34
p_{Col}	53	56	58	58	59	59
U_B	1800	595	248	133	83	19
U_C	1800	1600	1350	1152	1000	595
U_{Col}	1800	1800	1800	1800	1800	1800

Consider next Industry HoIn in Table 5. When marginal cost is increasing, the firms in the Bertrand Oligopoly earn positive profits. The joint profit of Bertrand duopolists is substantially below the monopolist's profit, but has not fallen all the way down to zero; further increases of the number of firms drive the joint profit down to zero at a declining pace. Note that the joint profit of Cournot duopolists exceeds the monopolist's profit; the cost savings obtained by spreading production over two firms outweigh the depressing effect of the additional firm on the price of the good. From two firms onward industry profit declines towards zero, at a slower pace than in the Bertrand Oligopoly. The behaviour of industry profit in the "collusive" oligopoly is quite different: it moves upwards, back to its level in the classic case, as the number of firms increases. As shown in Appendix B, the Pareto optimal set of this industry is concave to the origin. The centroid Col is its point where industry output and profit are largest: dividing total demand evenly among the firms is the most profitable arrangement of the industry. The firms wish to stay small, because in that way they avoid the adverse effects of increasing marginal cost. An increase of demand, for example through a drop of the value of d , is most profitably met by the entry of firms, not by the growth of existing firms.

The value of $t - d$ in Industry HeCon equals the value of c_2 in Industry HoIn, so that the industries have identical profit functions in quantity space. Therefore a number of rows of Table 8 are identical to the corresponding rows of Table 5, those for U_{Col} , U_C , Q_{Col} , Q_C and Q_Z , to be precise. The matching prices do differ, because the industries have different inverse demand functions. The Bertrand Equilibrium is at a different point in both quantity space and price space. As to the "collusive" oligopoly, just like the firms in Industry HoIn, the firms in Industry HeCon wish to stay small, but for a different reason: they want to avoid the adverse effects of (relatively) fast decreasing marginal revenue. An increase of demand is most profitably met by the entry of firms that bring new varieties to the market, not by producing more of existing varieties.

Industry HeIn combines increasing marginal cost with "fast" decreasing marginal revenue, which may be typical of many traditional industries. As we have seen, both changes from Industry HoCon affect the quantities in the same direction. The figures in Table 6 confirm that the deviations of the quantities from their HoCon-counterparts are similar to, and larger than with one of the changes separately. As a consequence, this observation applies to the profits of the "collusive" oligopoly, too.

The picture is completely different for Industry HoDe, which produces one homogeneous good using a technology with decreasing marginal cost (see Table 9). Bertrand oligopolists suffer losses as long as their number exceeds one. Merger increases the losses of the industry, unless all firms merge at once into one firm, which then starts acting as a monopolist and makes a large profit. A given number of Cournot oligopolists (more than one) perceive a stronger incentive to merge than in any other industry here considered. The firms in the "collusive" oligopoly, too, perceive an incentive to merge, albeit less strongly than the Cournot oligopolists do; the reason is that industry profit at the (exogenously fixed) initial number of firms is already maximal and exceeds by far the joint profit of the same number of Cournot oligopolists. As shown in Appendix B, the Pareto optimal set of this industry is convex to the origin. The centroid Col is its point where industry output and profit are smallest. They reach their maxima at any of the monopoly points, because decreasing marginal cost is exploited maximally by concentrating all production in one firm. An increase of demand, for example through a drop of the value of d , is most profitably met by increasing the output of this one firm.

Industry HeDe, with decreasing marginal cost and product differentiation, may be characteristic of many modern industries. The sum of $t - d$ and c_2 of the industry is zero: the profit functions in quantity space are identical to those in Industry HoCon. The remarks made for Industry HeCon on this matter apply here as well. In the "collusive" oligopoly, the joint profit is constant across the number of firms; once more, joint profit maximisation does not select a specific size distribution of firms.

The flawed treatment of the vector maximisation problems by noncooperative game theory results in distorted incentives for merger: in five of the six industries, the Bertrand oligopoly and the Cournot oligopoly yield other outcomes for the optimal number of firms than the "collusive" oligopoly does. In the author's opinion, the "collusive" oligopoly agrees best with a cursory view of the world. Naturally, this observation carries no weight in a dispute on a mathematical issue.

5. Some comments on laboratory experiments

Oligopoly models like the ones studied here are popular tools in the economics laboratory for testing theories of behaviour in situations with few, interacting participants. In contrast to what the use of the term "laboratory" suggests,

however, the proceedings in the economics laboratory differ fundamentally from those in the physics laboratory. Whereas in the physics laboratory the participants (for example, elementary particles) “know” the laws of nature and the experimenters are struggling to find out what the laws are, in the economics laboratory the experimenters have set the “laws of nature” and the participants (often undergraduate students) are struggling to find them out. Can one rationally expect beginners to grasp, within an hour or so, the mechanics of an artificial world that has been a source of confusion for almost two centuries?

It will surely help to give the beginners a head start by instructing them extensively, on the model and the means at their disposal to reach good decisions. A concern of the designer of the experiment is to supply the participants with adequate information without unveiling the solution. However, suggesting certain procedures for attacking the problem seems quite justified. The procedures need not be sophisticated; after all, Huck et al. (2004) have shown that, for the industries HoCon and HoIn, a simple noisy trial-and-error method always leads to the centroid of the Pareto optimal set. Meanwhile, the designer must beware of leading the endeavors of the participants in a particular direction. This aspect gains weight in connection with another one. Participants in experiments have often been, and in future experiments will be recruited from undergraduate students. Probably many of them will be economics students, with prior exposure to economic theory and the fallacy of Nash Equilibrium.

Even when prospective participants have received extensive instructions, it may be a good idea to familiarise them as monopolists with the model and the experimental setting. A bad performance of some participant as a monopolist puts into perspective the outcomes of her later plays of games.

The last, but not least important comment on current practice in the economics laboratory is this. In the physics lab, the experimenters create the conditions in which the phenomena predicted by their models are likely to occur. Experimenters in the economics lab, however, have frequently failed to do so. The first instruction that the participants in an oligopoly game receive goes often like this:

During the experiment you are not allowed to talk to other participants. If something is not clear, please raise your hand and one of us will help you.

Nash (1951) is to blame for this ban on communication. He suggested a new “solution concept” for a game, which would apply when the players of the game were unable to communicate and cooperate. Essentially, Nash replaced a simultaneous maximum problem with a set of conditional maximum problems by postulating that each player conditions on the *endogenous* actions of the other players; when applied to an oligopoly game, the Nash Equilibrium is the Bertrand Equilibrium (Cournot Equilibrium) when the firms use the prices (quantities) as instruments. Mathematically, the postulate amounts to ignoring the partial cross-derivatives of the profit functions when the first-order conditions for a solution are derived, *as if they are zero*.⁸ The resulting conditions are incompatible with the first-order approximations of the profit functions, for—when evaluated at a Nash Equilibrium—the cross-derivatives turn out to be *non-zero*. We have a prime example of a contradiction here; oddly enough, in economic theory and game theory this familiar outcome has not been recognised as the sign of a logical flaw. Regardless of this mathematical error, the ban on communication in experiments is not in keeping with the specification of the model. Variables corresponding to the actions of communication and cooperation are not present in the model, let alone constraints on such actions. Certain constraints may be lurking in the background, but in the model under empirical scrutiny their Lagrange multipliers are zero. Stated differently, communication and cooperation are free activities in the artificial world of the model. Therefore, the first instruction better be replaced by something like this:

The experiment consists of a number of plays of a game. During the experiment, you and the other players have access to a chatroom, where you may discuss any issues concerning the plays of the game. However, if you have questions concerning the experimental setting, please raise your hand and one of us will help you.

The criticism above is not to deny the interest of experiments in which the participants may not communicate. It merely stresses that the model does not apply to this situation, and that a Nash Equilibrium is not a mathematically consistent yardstick to judge the outcomes. How the impossibility of communication affects the outcomes is a question that surpasses the boundaries of the problem as it has been posed. Impediments for collusion must be modelled explicitly, for example along the lines suggested by Stigler (1964).

⁸See Appendix A.2 for a summary of issues concerning the Nash Equilibrium.

6. Concluding remarks

Ever since Cournot (1838), economists have built multi-agent models in the following way. First, they derived the first-order conditions of each agent's conditional maximum problem, treating only the agent's actions as endogenous and conditioning on the actions of the other agents. Next, they assembled all first-order conditions into one system of equations, adding some new (!) constraints (the market equilibrium conditions) if needed. Von Neumann and Morgenstern (1947), and De Finetti (1937a,b) before them, noted that the conditional maximum problems were interdependent and actually constituted a vector maximisation problem, or game. Shortly after, Nash (1951) postulated that agents unable to cooperate would in fact condition on the *endogenous* actions of their rivals; thus he disassembled the game into the old set of conditional maximum problems. However, the justification of the postulate relies on a particular interpretation of the problem. To see the fallacy of Nash Equilibrium, just concentrate on the mathematics and solve the game devoid of any interpretation. The generalised Cournot Oligopoly serves well as an example in an introduction to the method of vector maximisation, meanwhile exposing the flaw in noncooperative game theory.

Nash (1951) uses a fixed-point theorem to prove the existence of a Nash Equilibrium; its status of fixed point of some mapping gives the Nash Equilibrium an aura of stability. Here, the existence of a Nash Equilibrium is shown constructively, by solving a system of linear equations. Because the games are simple, it is easy to study the properties of a Nash Equilibrium. Its dependence on the arbitrary choice of instruments nullifies any claim of stability: The prices at the Cournot Equilibrium do not constitute a Nash Equilibrium in price space, so that (according to noncooperative game theory itself) the firms will be tempted to change their prices until they arrive at the Bertrand Equilibrium, where they observe that the implied quantities do not constitute a Nash Equilibrium in quantity space, so that . . . , and so on.

In a review of the experimental literature, Haan et al. (2006) find that *The ability to communicate among sellers has a strong and positive effect on the ability to collude*. The finding is good news for the proponents of the rationality postulate as the starting point of economic theory. In real life, there appears to be more coordination of actions than is compatible with noncooperative game theory. Because an experimental setting that allows of easy communication is a better approximation of many real-world situations than the alternative, the finding is consistent with this observation. Individually rational decision makers seem to understand well that in many situations they serve their private interests best by acting in unison with others.

'There is such a thing as being just plain wrong' (Richard Dawkins in *The Selfish Gene*), and that's what non-cooperative game theory is. Economics cannot hold on to a failed solution of a vector maximisation problem as its prediction of the outcome of rational behaviour in situations where actions are interdependent. The theory now known as *cooperative* game theory is the basis of the theory of rational decision making, unqualified by adjectives like "cooperative" versus "noncooperative," or "individual" versus "collective."

References

- Bertrand, J., 1883. Revue de la "Théorie mathématique de la richesse sociale" et des "Recherches sur les principes mathématiques de la théorie des richesses". *Journal des Savants* 67, 499–508.
- Cournot, A.A., 1838. *Recherches sur les principes mathématiques de la théorie des richesses*. Hachette, Paris.
- Cournot, A.A., 1927. *Researches into the Mathematical Principles of the Theory of Wealth* (translation by N.T. Bacon of Cournot, 1838). Second ed., Macmillan, New York.
- De Finetti, B., 1937a. Problemi di "optimum". *Giornale dell'Istituto Italiano degli Attuari* 8, 48–67.
- De Finetti, B., 1937b. Problemi di "optimum" vincolato. *Giornale dell'Istituto Italiano degli Attuari* 8, 112–26.
- Edgeworth, F.Y., 1881. *Mathematical Psychics: An Essay on the Application of Mathematics to the Moral Sciences*. C. Kegan Paul & Co., London.
- Haan, M.A., Schoonbeek, L., Winkel, B.M., 2006. Experimental results on collusion: The role of information and communication.
- Huck, S., Normann, H.T., Oechssler, J., 2004. Through trial and error to collusion. *International Economic Review* 45, 205–224.
- Kuhn, H.W., Tucker, A.W., 1950. Nonlinear programming, in: Neyman, J. (Ed.), *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability* (held in 1951). University of California Press, Berkeley, pp. 481–92.
- Nash, J., 1951. Non-cooperative games. *Annals of Mathematics* 54, 286–95.
- Nieuwenhuis, A., Unpublished results, 2017. Reconsidering Nash: the Nash equilibrium is inconsistent. URL: <https://dx.doi.org/10.13140/RG.2.2.29069.03043>. Paper presented at the 21st Annual ESHET Conference, Rationality in Economics, May 18–20, 2017, Antwerp.
- Nieuwenhuis, A., Unpublished results, 2018. Comparative statics for oligopoly: Flawless mathematics applied to a flawed result. URL: <https://dx.doi.org/10.13140/RG.2.2.22776.29443>.
- Stigler, G.J., 1964. A theory of oligopoly. *Journal of Political Economy* 72, 44–61.
- Von Neumann, J., Morgenstern, O., 1947. *Theory of Games and Economic Behavior*. Second ed., Princeton University Press, Princeton, New Jersey.

Appendix A. Vector maximisation and Nash Equilibrium

Appendix A.1. Vector maximisation

Equivalence of forms

Section 2 identifies the oligopoly game as a vector maximisation problem with equality constraints. The problem is an example of a game in extensive form. Unlike optimisation theory, noncooperative game theory cannot deal with games in extensive form. It does not need to, because it may first transform an extensive form into a normal form. The defining property of a normal form—the notion of Nash Equilibrium is inextricably bound to it—is that the arguments of the maximands may be varied independently of one another. It allows the game theorist to assign each argument to a different player as the one controlling it. Still, the assignment is just part of a particular interpretation of the problem, which does not affect its solution. But the game theorist cannot ignore the issue of equivalence of the various forms that arises in the transformation: all forms must yield the same outcome, else the transformation is not permissible. Non-singularity of the transformation guarantees equivalence of the forms.

A first-order condition

The mathematician deals with the Cournot Oligopoly by considering the first-order approximation $\mathbf{U}d\mathbf{q} = d\mathbf{v}$, where $\mathbf{U} := [\partial u_i / \partial q_j]$. A necessary condition for the outcome to be (weakly or strongly) Pareto optimal at a point \mathbf{q}^* is that the matrix $\mathbf{U}(\mathbf{q}^*)$ have deficient row rank (else a solution of $\mathbf{U}d\mathbf{q} = d\mathbf{v}$ with a strictly positive or positive vector $d\mathbf{q}$ exists such that $\mathbf{U}d\mathbf{q} = \mathbf{0}$). Therefore the Pareto Optimum is (a part of) the variety on which $|\mathbf{U}| = 0$. At a Pareto optimal point a vector $d\mathbf{q}$ exists such that $\mathbf{U}d\mathbf{q} = \mathbf{0}$. The geometric interpretation is that the iso-profit curves have a tangent line in common at points of the Pareto Optimum. The other part of the first-order conditions is not considered here, nor are the second-order conditions.

The derivation varies all quantities simultaneously, and *it varies every quantity in all places where it occurs*. Not doing so would be a grave mathematical mistake.

In the case of a single maximand, the condition that the $(1 \times I)$ -dimensional matrix, or row vector of partial first-order derivatives have deficient rank amounts to the condition that all first-order derivatives vanish. The condition that derivatives parallel to the coordinate axes be zero does *not* move over from scalar to vector maximisation.

Invariance of the solution

Let $v_i(p_1, \dots, p_I), i = 1, \dots, I$, be the profit functions in price space. Define $\mathbf{V} := [\partial v_i / \partial p_j]$, $\mathbf{F} := [\partial f_i / \partial q_j]$, and $\mathbf{G} := [\partial g_i / \partial p_j]$. Application of the chain rule yields $\mathbf{U} = \mathbf{V}\mathbf{F}$ and $\mathbf{V} = \mathbf{U}\mathbf{G}$. Because $|\mathbf{U}| = |\mathbf{V}| \cdot |\mathbf{F}|$ and $|\mathbf{F}| \neq 0$ (or, because $|\mathbf{V}| = |\mathbf{U}| \cdot |\mathbf{G}|$ and $|\mathbf{G}| \neq 0$), $|\mathbf{U}| = 0 \iff |\mathbf{V}| = 0$: the quantity-setting oligopoly and price-setting oligopoly yield the same outcomes, because tangency of iso-payoff surfaces is preserved. In fact, the solution of every optimisation problem is invariant under nonsingular transformations of variables.

“Conjectural variations”

Take a point where $\mathbf{U}d\mathbf{q} = \mathbf{0}$, and let \mathbf{U} have rank $I - 1$. Then $d\mathbf{q}$ is unique up to a factor and proportional to the vector of cofactors $|\mathbf{U}_{ij}|^*$ of the elements of a row of \mathbf{U} , for $\sum_j u_{ij} |\mathbf{U}_{ij}|^* = |\mathbf{U}| = 0, i = 1, \dots, I$. Define $\xi_{ji} := dq_j / dq_i = |\mathbf{U}_{ij}|^* / |\mathbf{U}_{ii}|^*$, the marginal rate of substitution between q_j and q_i along the common tangent line to the level curves of the maximands at the point. Obviously, $\xi_{jk} \xi_{ki} = \xi_{ji}, i, j, k = 1, \dots, I$. The matrix $\Xi := [\xi_{ji}]$ has rank 1, and all its diagonal elements are 1: each column $\xi_{\cdot i}$ of Ξ is a differently scaled version of the vector in the nullspace of \mathbf{U} .

In oligopoly theory the entities ξ_{ji} are known under the misnomer of “conjectural variations.” They are treated as free parameters, which may be given alternative values to represent different ways of behaviour. In fact, their values are determined endogenously, as part of the solution of the vector maximisation problem. The first-order conditions of a “conjectural variations” oligopoly are just the condition $\mathbf{U}d\mathbf{q} = \mathbf{0}$ in disguise.⁹

In a *symmetric* oligopoly model like the one studied here, $\mathbf{U}\mathbf{1} = \mathbf{0}$ (where $\mathbf{1}' := [1 \dots 1]'$) at *Col* and at no other point (the directional derivatives of the profit functions along the symmetry axis are zero at *Col*). This outcome explains why the thought experiment on Page 4 of varying the quantities in unison (*i.e.*, $d\mathbf{q} = \varepsilon\mathbf{1}$; all “conjectural variations” equal to 1) leads to *Col*.

⁹See Nieuwenhuis (Unpublished results, 2017, Sections 2.2.3, 2.3.4 and 4.1.5) for more on “conjectural variations.”

Appendix A.2. Nash Equilibrium

Definition

Inspired by the economic interpretation of the problem ('The only thing that a firm gets to control is its own output, not the output of other firms.'), noncooperative game theory postulates $u_{ii} = 0, i = 1, \dots, I$, as the first-order conditions for maximal profits of the firms. At a Cournot Equilibrium, the iso-profit curves are *perpendicular* instead of tangent.

Please note that, against the very nature of mathematics, the classic vindication of the Nash Equilibrium relies on a specific interpretation of the vector maximisation problem. In fact, the interpretation has just led the game theorist astray. The inadequacy of the reasoning may be argued in several ways.

Internal inconsistency

As noted on Page 10, the cross-derivatives of the profit functions, implicitly set to zero in the definition of Nash Equilibrium, turn out nonzero at the Cournot Equilibrium. Here is a prime example of a contradiction.

Flawed derivation

The approach of the mathematician adds to the reasoning of the game theorist. A firm will reflect on how alternative levels of its output affect its profit. A change of its output level affects the profits of its competitors as well, and its competitors, too, will reflect on changes in their output levels that affect its own profit. In other words, although each firm has control only over its own output, it still has to endure the changes in the outputs controlled by the other firms. In this way all firms arrive at a problem identical to that of the mathematician. Starting from some position, the firms may go looking for win-win moves until no such move is left.

Clearly, that each firm controls only its own output is not a valid argument for neglecting the cross-derivatives of the profit functions. In fact, varying a variable at one place while holding it constant at other places is just a grave mathematical mistake, for which no excuse exists at all.

Another sophism in the story telling that surrounds the Nash Equilibrium is that the players condition on the actions of their rivals because they cannot communicate and cooperate. Again, mathematics comes to the rescue and dismantles it. As was pointed out in Section 5, communication and cooperation are free activities in the model because their shadow prices are zero.

The predominance of noncooperative game theory justifies a variation on Edgeworth (1881, p. 127, fn. 3): *To treat variables as constants is the characteristic vice of the ~~un~~mathematical economist.*

Faulty outcome

A Nash Equilibrium is not Pareto optimal. The domains of the arguments contain points where all profit functions have larger values. In the terms of game theory: the firms may all increase their profits by moving to some other point in their action possibility sets, for the model contains no restrictions that keep them from doing so. Noncooperative game theory does not describe profit maximising behaviour. What motive, then, guides the behaviour of the firms?

Because they are mathematically equivalent forms of a vector maximisation problem, the quantity-setting oligopoly and the price-setting oligopoly must yield the same outcomes. As shown above, they do when treated correctly. Noncooperative game theory, however, results in the Cournot Equilibrium and Bertrand Equilibrium, which are not identical; unlike tangency, perpendicularity of iso-payoff surfaces is not preserved in all nonsingular transformations.

May we one day welcome a new "solution concept" for the problem of budget-constrained utility maximisation, one that yields a different outcome for every other choice of argument to be eliminated from the utility function?

(Lack of) stability

The classic defence of the notion of Nash Equilibrium is that it would be stable, or self-enforcing: a Nash Equilibrium is a point at which no firm has an incentive to deviate *given what other firms are doing*. As pointed out on Page 11, it fails because a Nash Equilibrium depends on the frame of reference. The prices at the Cournot Equilibrium do not constitute a Nash Equilibrium in price space, nor do the quantities at the Bertrand Equilibrium constitute a Nash Equilibrium in quantity space. According to noncooperative game theory itself, at a Cournot Equilibrium each firm would regret its choice of price and want to change it unilaterally, and similarly at a Bertrand Equilibrium the firm would regret its choice of quantity and want to change it unilaterally.

The classic attack of noncooperative game theory on the Collusion Point is that it would *not* be self-enforcing: if firms found themselves there, then each would have the incentive to *unilaterally* deviate from that output level. This is certainly correct, but it is a feature, not a bug of the notion of Pareto optimality. As noted in Appendix A.1, the condition that derivatives parallel to the coordinate axes be zero does not move over from scalar to vector maximisation. In fact, just like the classic defence of the Nash Equilibrium fails, so does the classic attack on the Collusion Point. In both cases, the argument relies on a specific interpretation of the mathematical problem and involves unwarranted conditioning on endogenous variables.

To see the odd behaviour that noncooperative game theory imposes on the firms, put quantity-setting firms of Industry HoIn at *Col*, which is the industry's Joint Profit Maximum, and see where it leads them to. When all firms yield to the incentive of deviation, they go down Profit Hill. They persevere in the direction that lowers their profits till they reach the Cournot Equilibrium. Price-setting firms will, after arrival at the Cournot Equilibrium, crash even further down to the Bertrand Equilibrium (so much for the Cournot Equilibrium being self-enforcing). The simple trial-and-error method of Huck et al. (2004), however, lets all firms change their output levels by a small amount. The firms that, as the outcome of the simultaneous move, notice a fall of their profits make a U-turn in the next step, and so on. Starting from any point, for example a Nash Equilibrium, the firms arrive at *Col*. Which approach describes rational, optimising behaviour?

A stepwise method

In scalar optimisation problems, the method of stepwise maximisation may be convenient; examples are concentrating the likelihood function, and cost minimisation conditional on some unspecified output level, which is yet to be determined in a subsequent step of profit maximisation. As noted on Page 3, however, the “first-order conditions” for a Nash Equilibrium are *not* the first step in such a procedure, as there is no way of undoing the conditioning at a subsequent step. Still, there is a stepwise method also in vector maximisation problems. Take a duopoly game as an example. In the first step, Firm 1 picks one of its iso-profit curves and allows Firm 2 to choose the point on the curve where its profit is maximal; in subsequent steps, the firms repeat this exercise for all permissible values of the first firm's profit. They may exchange their roles just to find that they arrive at the same solution, the Pareto Optimum. The procedure works because each firm, rather than condition on what the rival firm does, conditions on what it wants. Like in a detective, the crucial issue is motive.

Appendix B. The solution of the oligopoly game

A change of the units of measurement simplifies the formulas of model and solution. Define the scaled quantities $x_i := q_i/k$, $i = 1, \dots, I$, $k := (d_0 - c_1)/d$, and divide the profit functions (5) by dk^2 to redefine them to

$$\begin{aligned} v_i &= u_i(x_1, \dots, x_I) = (1 - X)x_i - rx_i^2, & i = 1, \dots, I, \\ r &:= \frac{t - d + c_2}{d}, \end{aligned} \tag{B.1}$$

where $X := \sum_j x_j$. The expressions for x_{Col} and v_{Col} are

$$x_{\text{Col}} = \frac{1}{2(I + r)}, \tag{B.2a}$$

$$v_{\text{Col}} = (I + r)x_{\text{Col}}^2. \tag{B.2b}$$

According to (B.2), $-1 < r$ must hold for optimal monopoly output and profit to be positive and finite. The expression for p_{Col} may be written as

$$p_{\text{Col}} = \frac{(I + r - s)c_1 + (I + r + s)d_0}{2(I + r)},$$

where $s := c_2/d$. The larger is c_2 , the smaller is the weight of c_1 and the larger is the weight of d_0 , and hence the higher is the price at the Collusion Point.

Appendix B.1. The maximal-profit set

Consider first the equivalent of the fact stated in Section 2, p. 4,

FACT: The maximal-profit set of oligopoly (B.1) is the set of points where the industry's profit-output ratio equals 1/2.

PROOF: The proof is for a duopoly, but the method of proof applies to any number of firms. Consider the industry's profit along the ray $x_2 = \alpha x_1$, $\alpha \geq 0$. Substitute $x_2 = \alpha x_1$ in the profit functions of the firms and add the outcomes:

$$v_1 = v_1(x_1, \alpha) = x_1 - (1 + \alpha + r)x_1^2, \quad (\text{B.3a})$$

$$v_2 = v_2(x_1, \alpha) = \alpha x_1 - (\alpha + \alpha^2(1 + r))x_1^2, \quad (\text{B.3b})$$

$$\Upsilon = V(x_1, \alpha) = (1 + \alpha)x_1 - ((1 + \alpha)^2 + (1 + \alpha^2)r)x_1^2. \quad (\text{B.3c})$$

Let Υ_α be the maximal Υ along the ray $x_2 = \alpha x_1$, reached when $x_1 = x_{1,\alpha}$. The ratio $\Upsilon_\alpha/x_{1,\alpha}$, is $(1 + \alpha)/2$, so that the profit-output ratio of the industry is in fact 1/2. Because the profit-output ratio is strictly decreasing in industry output, there is no other point along the ray where the profit-output ratio of the industry is 1/2. *Q.E.D.*

To gain a better understanding of the outcome, rewrite (B.3b) and (B.3c) to

$$w_2(x_1, \alpha) = x_1 - (1 + \alpha(1 + r))x_1^2, \quad \alpha > 0, \quad (\text{B.3b}')$$

$$W(x_1, \alpha) = x_1 - \left(1 + \alpha + \frac{1 + \alpha^2}{1 + \alpha}r\right)x_1^2. \quad (\text{B.3c}')$$

(B.3a), (B.3b') and (B.3c') are quadratic functions of the form $ax^2 + bx$ with a common value of b but different values of a . The value of a in the third function is a weighted average of its values in the first two functions, with the weights of $1/(1 + \alpha)$ and $\alpha/(1 + \alpha)$, respectively. Therefore the point on the ray where industry profit is maximal is generally in between the points where the profits of the firms reach their maxima. As the ray $x_2 = \alpha x_1$ revolves around the origin from the x_1 -axis ($\alpha = 0$) towards the x_2 -axis ($\alpha = \infty$), the weight of Firm 1's profit function declines from 1 to 0. Simultaneously, the intersection of the ray with the zero-profit line of Firm 1, $Z_{1,\alpha}$, moves from the point $Z_1 = Z_{1,0} := [1/(1 + r), 0]$ towards the point $B_1 = Z_{1,\infty} := [0, 1]$. Parallel to this line (segment), at half the distance from the origin, is the maximal-profit line of Firm 1; along the line, profit declines linearly from the monopoly profit at $C_1 = P_{1,0} := [1/(2(1 + r)), 0]$ towards zero at $P_{1,\infty} := [0, 1/2]$. The maximal-profit line of Firm 2 is the mirror image of the one of Firm 1 with respect to the ray $x_1 = x_2$.

When $r = 0$, the maximal-profit lines of both firms are $1 - 2X = 0$, which is the maximal-profit set and Pareto optimal set. Along the line, industry output and profit are constant. One may show that, at every point of the line, the ray $x_2 = \alpha x_1$ through the point is the common tangent line of the iso-profit curves of the firms.

In the general case of $r \neq 0$, the maximal-profit lines intersect at *Col*. The industry has a maximal-profit curve, which runs from C_1 through *Col* to C_2 , the mirror image of C_1 with respect to $x_1 = x_2$; the intermediate segments of the curve are in between the maximal-profit lines. The ray $x_2 = \alpha x_1$ is tangent to an iso-profit curve of the industry where it intersects the maximal-profit curve. The iso-profit curves of the firms do *not* have a common tangent line at the intersection, unless it is at *Col* or a monopoly point.

How do industry output and profit evolve along the maximal-profit curve? Compare the output of a firm at *Col*,

$$x_{\text{Col}} = \frac{1}{2(2 + r)}, \quad (\text{B.2a}')$$

to the output of a firm halfway between C_1 and C_2 , which equals half the monopolist's output,

$$\frac{x_{\text{mono}}}{2} = \frac{1}{4(1 + r)}.$$

If $r > 0$, then $x_{\text{Col}} > x_{\text{mono}}/2$, so that the curve is concave to the origin; it proves to be a segment of an ellipse (see below). *Col* is the point of the curve where industry output and profit are largest. If $r < 0$, then $x_{\text{Col}} < x_{\text{mono}}/2$, so that

the curve is convex to the origin; it proves to be a segment of a hyperbola (see below). Col is the point of the curve where industry output and profit are smallest, they are largest at any of the monopoly points.

For any number of firms, putting the sum of the profit functions equal to $X/2$ yields the equation of the maximal-profit (hyper)surface,

$$-2X^2 - 2r \sum_j x_j^2 + X = 0. \quad (B.4)$$

If $r = 0$, the quadric surface reduces to $X(1 - 2X) = 0$. Because $X \neq 0$, the solution is $1 - 2X = 0$ (which in this case is the Pareto optimal set). For a duopoly, the Discriminant Δ of the cone section is $-16r(2+r)$. If $r > 0$ (and if $r < -2$), then $\Delta < 0$: the equation represents an ellipse. If $-2 < r < 0$, then $\Delta > 0$: the equation represents a hyperbola. For larger numbers of firms, the equation represents ellipsoids or hyperboloids of revolution.

Appendix B.2. The Pareto optimal set

Turn next to the Pareto optimal set. When there are I firms, $|\mathbf{U}|$ is a polynomial of degree I in I variables. If $r = 0$, the condition $|\mathbf{U}| = 0$ reduces to $(1 - X)^{I-1}(1 - 2X) = 0$, which represents the Pareto optimal set and Zero-profit set; the Pareto optimal set and maximal-profit set coincide. For a duopoly, the Discriminant Δ of the cone section is $16(1+r)^2r(2+r)$. If $r > 0$ (and if $r < -2$), then $\Delta > 0$: the equation represents a hyperbola. If $-2 < r < 0$, then $\Delta < 0$: the equation represents an ellipse. The Pareto optimal set and maximal-profit set do not coincide. They do have the Collusion Points of the included oligopolies in common; at the intermediate segments the maximal-profit set is slightly farther away from the origin than the Pareto optimal set. The property of a constant profit-output ratio does not apply to the Pareto optimal set. Still, the evolutions of industry profit and output along the Pareto optimal curve and maximal-profit curve are qualitatively similar. Figure B.1 sketches the solution of $|\mathbf{U}| = 0$ for the duopoly HoIn.

What does the Pareto optimal set look like when the number of firms exceeds two? In a triopoly, there are three monopoly points and three curves like C_1C_2 connecting them. The Pareto optimal set is a surface area, topologically a triangle, the three sides of which are the curves of the included duopolies. The surface area is convex to the origin, flat, or concave to the origin for $r < 0$, $r = 0$, or $r > 0$, respectively. For a tetrapoly, the Pareto optimal set is a volume, topologically a tetrahedron, the four faces of which are the “triangles” of the included triopolies. And so on, beyond graphical representation, for still larger numbers of firms. Here is a prime example of a result, for the first time stated and proved by De Finetti (1937a, Section 12; translation of the author),

The locus of “optimum” points with respect to n functions is, topologically, a simplex of $n - 1$ dimensions, the n faces of which are the loci of “optimum” with respect to $n - 1$ <of the> functions, the $\binom{n}{2}$ edges of which those for $n - 2$ <of the> functions, and so on, up to the n vertices, “optimum” points with respect to the n functions separately.

The “locus of “optimum” points” is what we call the Pareto optimal set nowadays.




Figure B.1: Industry HoIn ($r = c_2/d = 1/4$)

Notes: The curve through B_1 , B and B_2 solves $|U| = 0$, but is not part of the Pareto Optimum. The curve through C_1 and C_2 solves $|U| = 0$, and is the Pareto Optimum; its midpoint is the Collusion Point. $[B_i Z_i]$ is the Zero-profit line of Firm i . Z is the Zero-profit point, B'_i the end point of Firm i 's Bertrand reaction function, C_i the end point of Firm i 's Cournot reaction function (also Firm i 's monopoly point). At the Bertrand Equilibrium B the iso-profit curves have a tangent line in common, but still they intersect. At the Cournot Equilibrium C the iso-profit curves are perpendicular to one another.

Source: Nieuwenhuis (Unpublished results, 2017, Figure 4).