



On Dynamic Games with Randomly Arriving Players

Pierre Bernhard, Marc Deschamps

► To cite this version:

Pierre Bernhard, Marc Deschamps. On Dynamic Games with Randomly Arriving Players. Dynamic Games and Applications, 2016, pp.34. 10.1007/s13235-016-0197-z . hal-01425013

HAL Id: hal-01425013

<https://inria.hal.science/hal-01425013>

Submitted on 3 Jan 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

On dynamic games with randomly arriving players

Pierre Bernhard* and Marc Deschamps†

September 24, 2015

Abstract

We consider a dynamic game where additional players (assumed identical, even if there will be a mild departure from that hypothesis) join the game randomly according to a Bernoulli process. The problem solved here is that of computing their expected payoff as a function of time and the number of players present when they arrive, if the strategies are given. We consider both a finite horizon game and an infinite horizon, discounted game. As illustrations, we discuss some examples relating to oligopoly theory (Cournot, Stackelberg, cartel).

Keywords Dynamic game, Bernoulli process of entry, Oligopoly

JEL classification : C72, C61, D21, L13

*BIOCORE team, INRIA Sophia Antipolis-Méditerranée, France. *Pierre.Bernhard@inria.fr*

†Université de Franche-Comté CRESE, BETA-CNRS and OFCE-Sciences Po.

Marc.Deschamps@univ-fcomte.fr

We thank Céline Savard-Chambard, Saïd Souam, Nicolas Eber and Hervé Moulin for conversations and comments that improved the paper. We received very useful suggestions from Sylvain Béal. Marc Deschamps acknowledges the financial support of ANR Damage (programme ANR-12-JSH1-0001, 2012-2015). The usual disclaimer applies.

1 Introduction

In most strategic interactions analysis the set of players is known and common knowledge. But there are some situations where neither the players, nor the designer or public authority, nor the theorist know how many players *are* in a game. As an illustration, consider the three situations below:

Example 1

Imagine the case of a company that commits a violation of competition rules (which identically affects consumers) in a country where consumers can take joint action in court to recover the amount of their losses. The legal system chosen by the country as part of a joint action works as follows: in the first period a consumer enters the court, in the second period the judge analyses the case. If he finds it justified, he evaluates the total harm done¹, and imposes a reimbursement and a payments schedule for each month of a third period. But during that third period, more complainants may apply to join the action, given that the judge will not approve more than one per month. The monthly reimbursement is then shared equally among all complainants approved at that time. What is the first consumer-victim expected payoff according to the fact that he does not know how many victims will participate in the joint action as each victim can alternatively choose to never suit or suit individually the company?

Example 2

Consider the case of a city where the authorities issue a number of permits to individuals so that they can transport by car, for a fee, people in the city. There are k individuals who have such an authorization. Then, thanks to technological innovation, a company is able to transform every willing motorist driver into a transporter of people for fee. How much individuals k who have the permission of the city can expect, knowing that they do not know how many motorists will be turned into drivers?

¹On the one hand, in a joint action, this country chooses to charge the total harm to the liable party. On the other hand, it allows each victim not to prosecute, prosecute individually or collectively. It thus follows that the liable party can be led to pay several times the damage which she caused: these are punitive damages.

Example 3

Suppose that a French rich entrepreneur decides to create the “René and Gisèle Foundation” to promote the research in game theory. This foundation, thanks to the proceeds of a charitable donation, appoints a benevolent research director (free of conflicts of interest) to share each month equally a sum of 100.000 euros between all scientific projects in game theory which she has examined and found valuable. As the entrepreneur and the research director want to provide a long term support, they decide to never abandon the funding of a project that she had already selected. After the beginning of the foundation, where $k \geq 1$ projects could have been selected, each month, due to administrative procedures, only one new project could be added to the foundation’s public list of selected projects (which mentions the date of arrival in the list of each selected project). How much the m – *th* selected project can expect to have?

The goal of this paper is to provide a general mathematical framework for computing the payoff of players in a dynamic game with clone players randomly arriving at each step of time in a Bernoulli process. The paper is organized as follows: in the next section we expose the theoretical related literature and explain how we contribute to the question of games with an unknown number of players. In section 3, we formulate the problem, propose a model and give the expected payoff of finite (*Theorem 1*) and infinite (*Theorem 2*) horizon game, that is the expected payoff for each player. Then in Section 4 we provide some examples from oligopoly theory where this framework and results could be used to find expected payoff and static equilibria. Section 5 concludes by underlining the limits and questions of our framework which are left for future research.

2 Related literature

It is well known that the usual-textbook game theory toolbox is in the “fixed- n ” paradigm and so provides directly no solution to problems where there is a degree of uncertainty on the number of players in a game. Indeed, in Bayesian games, the uncertainty is on the *type* of each other player’s, that is on their private information about their characteristic represented by a random variable for each player. So there is no uncertainty on the number of players in the game, the uncertainty is on their characteristics. Moreover to extend Bayesian games to games with unknown number of players raises conceptual and modelling issues that could not be overcome. However, to the best of our knowledge, there exists currently in game theory

two ways to directly manage the question of uncertainty on the number of players in a game.

The first strain of literature, which is the oldest one, is the study of games where the *active competing players* is not common knowledge. This issue has taken two forms (see [Levin and Ozdenoren, 2004]). The first form models the number of active players with a stochastic process on a set of potential players. It means that the active players is a subset of the common knowledge potential players, which implies that each player has a private information (he is active or not). The timing of this kind of game is usually in four-step: 1/ nature chooses which players will be active (and whose those will be passive), 2/ each player observes only his own situation, 3/ the active players choose their strategies, and 4/ payoffs are received. The common knowledge information is the set of potential players and the probability distribution used by the nature to assign a role to each player. The second form endogenizes the entry process. In this setting, at first, there is *no player* in the market. Then, there is the assumption that each potential player (the total set of which is common knowledge) has a privately known and different cost to entry on the market, and each of them must decide to participate or not into the game. These two kinds of models are useful in a great number of cases. For example, consider the situation of an appointment of Dean at a University, the applicants do not know how many persons *will* apply even if they perfectly know how many persons *can* apply (the set of Professors in the University). It is also the case, for example, in online auction sites where the number of registers is known and the number of bidders for each product is *ex ante* unknown.

Seminal articles in this line of literature are [McAfee and McMillan, 1987a] and [Matthews, 1987] which demonstrate that some results of auction theory could be sensitive to the assumption that there is a stochastic number of bidders, which means that the set of real bidders is not common knowledge whereas the set of potential bidders is. They also allow the bidders to have different priors on how many bidders are present as long as these priors are Bayesian consistent. From the design policy point of view these articles open the question whether the seller had an interest to reveal (or conceal) the number of actual bidders and their identities. In [Piccione and Tan, 1996] the number of potential bidders is common knowledge but some of them are experts (bidders with private information) and some are non-experts (bidders with no private information), and it is the number of experts which is uncertain for the seller and the other buyers. [Münster, 2006] offers a study of rent-seeking contests with an unknown number of competing contestants where all players are risk neutral or have constant absolute risk aversion. Complementing to the pure risk approach, [Salo and Weber, 2007] and [Levin and Ozdenoren, 2004] propose models where bidders are adverse to the ambiguity generated by the lack

of information about the number of rivals in the auction. On the entry process endogenous side, [Samuelson, 1985] finds that policies to limit the number of bidders may be welfare improving. When a first-price sealed bid auction is used and bidders can enter upon paying an entry cost, [McAfee and McMillan, 1987b] show that the seller should not impose a reserve price higher than his own valuation and that the optimal number of bidders enter. More recently, there has been investigations on environmental topics as in [Rubio and Ulph, 2007] or [Breton et al., 2010, Breton and Keoula, 2012].

The more recent other strain of literature is games with population uncertainty, where the *real number* of players in the market is not common knowledge. But the number of players really in the game is supposed to be drawn from a random variable, whose probability distribution and mean are commonly known. Players have to choose their strategies before the player set is revealed. Among these games, a special attention has been served to the sub-class called *Poisson games*, where the number of players N is a random variable that follows a Poisson distribution with mean n (so $N = k$ with the probability $\frac{e^{-n}n^k}{k!}$), and each player's type is then drawn independently from some fixed probability distribution. More generally, in *extended Poisson games* the expected population sizes and players' utility functions may depend on an unknown state of the world. In this case there is a two-stage structure: first a random state variable is drawn from some given distribution, and then a Poisson game is played. In these games, each player's belief about the number of other players and their types is the same with the prior distribution of the total number of players and their types (property called by Myerson *environmental equivalence*). It has been shown that Poisson games modelling has very nice mathematical features, subsume Bayesian games with consistent priors and open a fertile field of research. To illustrate the nature of questions solved by it [De Sinopoli and Pimienta, 2009] introduces the following simple and clear story: "A player is sitting at home and faces two possible alternatives, either she goes out to some social event, or she stays home. She does not know how many players are facing this same disjunctive, but she knows that this number is a Poisson random variable with parameter n . If she goes out and meets somebody she receives a payoff equal to 1. If she meets nobody or decides to stay home, she gets a payoff equal to 0. Every player faces these same two options and has the same preferences".

This type of modelling was introduced by [Myerson, 1998b], [Myerson, 1998a], [Myerson, 2000] and [Milchtaich, 2004]. See also [De Sinopoli et al., 2014] for a more recent account. It seems particularly well adapted to elections (where each voter does not know what is the real number of voters, see Myerson's works) or some kind of auctions. [Makris, 2008] proposes to model the coordination problem

as a Poisson game and investigates the conditions under which a unique equilibrium is selected. In [Ritzberger, 2009] it is shown that with population uncertainty, two competitors are not enough to eliminate all profits in equilibrium and have competition in a Bertrand game. Recently, [Östling et al., 2011] proposes an application of Poisson game to a LUPI game (lowest unique positive integer wins lottery), in order to study a lottery called Limbo introduced by the government-owned Swedish gambling monopoly Svenska Spel on January 29, 2007.

By contrast with these two lines of literature, the spirit of our paper is a dynamic framework where the number of players involved varies with time, and to study the expected payoff of each player in situations where nobody knows *ex ante* the set of real players at each period after the beginning and at the end of the game (finite case), or at each period of the game (infinite case). The only information that are common knowledge before the game actually starts in our setup are the size of the market, the number of incumbents and that there is a random entry at each period of time according to a Bernoulli process of known parameter. Once the game is going, the actual number of players present at each stage is also common knowledge at that stage. So there is no asymmetric information in our model. We will consider a *clone economy*, i.e an economy where the incumbents and the new entrants are perfectly identical. Even though there is an equilibrium at each step, our aim at this stage will not be to define a dynamical equilibrium, but to compute the expected payoff to the players when their strategic behaviors are given. Therefore we have a sequence of static equilibria, but not a dynamic equilibrium. (See nevertheless appendix D).

In the rest of the literature on games, but with a fixed number of players, we may find some kinship of our theory with repeated games, since the folk theorem, for instance, involves an infinite sequence of plays with the same strategies, as we do. But there, it is the same set of players who play again and again. The fact that in the present theory, the payment matrices vary from step to step is reminiscent of stochastic games, see *e.g.* [Neyman and Sorin, 2003] or, more recently, the special issue [Nowak et al., 2013]. However, we are not, at this stage, able to prove a result about a stationary dynamic equilibrium, as often found in that strain of literature.

Finally, the literature the closest to our theory that we are aware of may be that of piecewise deterministic systems [Davis, 1985, Haurie et al., 1994], although this is not in a game theoretic framework. But we may consider the stochastic arrivals of players as stochastic jumps in the dynamic system, which is deterministic between two jumps. We show in appendix D how that theory could be put to use to compute a dynamic equilibrium. But at this stage, applying these ideas to concrete examples seems out of reach.

3 Model

The family of problems we are investigating could have two forms. As in our examples 1 and 3, there are cases where the dynamic game has an end which corresponds to a finite horizon. There are also cases, as in our example 2, where the dynamic game has no finite end-time or, more realistically, where the end-time is unknown for the players, which corresponds to the infinite horizon case. We will study these two types of games successively.

We consider a discrete time, multistage game where time t is an integer varying from t_1 to T (finite horizon), or from 0 to ∞ (infinite horizon). We assume that additional players may arrive only one at each discrete instant of time, this with a probability p and independently from other arrivals (a Bernoulli process), and let t_m be the first stage where there are m players, i. e. the stage when player number m arrives.

Usually, in a dynamic game there is a state $x(t) \in \mathbb{R}^n$ whose dynamics are driven by the actions $a_i(\cdot)$ of the players, according to some given law

$$x(t+1) = f(t, x(t), a_1(t), a_2(t), \dots, a_n(t)), \quad x(t_1) = x_1.$$

Then the payoff of each player is

$$\Pi_i = \sum_{t=t_1}^T L_i(t, x(t), a_1(t), a_2(t), \dots, a_n(t)).$$

If the a_i are mixed strategies, then L_i is a mathematical expectation.

However, here the emphasis is not on the strategies but on computing an expected payoff when the strategies are assumed identical for all players, and known as a function of the number m of players present, the state, and current time. In that case, a sufficient description of the state is the sequence $\{t_n\}_{n \leq m}$ of past arrival times, and current time t . As a consequence, we may dispense with the explicit dynamics, and write the payoff as a function of these variables. We also assume that the players are identical not only in their strategy choices, but also in their per stage payoff. (The latter justifies the former. We will however see a marginal deviation from that hypothesis in our examples.) Their payoffs will only differ because of their different arrival times.

Notation Let $\tau_1 = t_1$, and for $m \geq 2$, $\tau_m = (t_2, \dots, t_m) \in \mathbb{N}^{m-1}$, and therefore $\tau_{m+1} = (\tau_m, t_{m+1})$. Let $L_m(\tau_m, t)$ be the per stage payoff of each player when m players are present. Let also $m(t)$ be the number of players present at time t .

We introduce the further notation $M_1(t) = L_1(t)$, and, for $m \geq 2$,

$$\mathcal{T}_m(t) = \{\tau_m \mid t_1 < t_2 < \dots < t_m \leq t\}, \quad M_m(t) = \sum_{\tau_m \in \mathcal{T}_m(t)} L_m(\tau_m, t).$$

(in the infinite horizon case, we will use $t_1 = 0$). For example,

$$M_2(t) = \sum_{t_2=t_1+1}^t L_2(t_2, t), \quad M_3(t) = \sum_{t_2=t_1+1}^{t-1} \sum_{t_3=t_2+1}^t L_3(t_2, t_3, t).$$

We may provide a graphical representation of what is going on. All possible entrance histories, for a finite horizon, can be represented on a tree as in figure 1. Time is shown on the bottom line. At each instant of time, and at each node of the tree for that instant, we draw a branch going up if a new player enters at that time instant, going down if nobody enters. We have chosen $t_1 = 0$, shown four time steps, and labelled the nodes with the number of players present upon reaching this node (during the time step ending there). We have noted over each branch the per stage payoff of each player during the corresponding time step.

An history is represented by a path from the root of the tree (to the left) to a leaf. The payoff for that history is the (“horizontal”) sum of the per stage payoffs noted on each branch of its path (τ_m is denoted with a list of time instants between curly braces). For each m and t , $M_m(t)$ is the (“vertical”) sum of the L_m of stage t , i.e. the sum of the markings of the branches of the stage beginning at t reaching a node labelled m .

One can read on this graph that:

$$\begin{aligned} M_1(0) &= L_1(0), \\ M_1(1) &= L_1(1), \\ M_2(1) &= L_2(\{1\}, 1), \\ M_1(2) &= L_1(2), \\ M_2(2) &= L_2(\{1\}, 2) + L_2(\{2\}, 2), \\ M_3(2) &= L_3(\{1, 2\}, 2), \\ M_1(3) &= L_1(3), \\ M_2(3) &= L_2(\{1\}, 3) + L_2(\{2\}, 3) + L_2(\{3\}, 3), \\ M_3(3) &= L_3(\{1, 2\}, 3) + L_3(\{1, 3\}, 3) + L_3(\{2, 3\}, 3), \\ M_4(3) &= L_4(\{1, 2, 3\}, 3). \end{aligned}$$

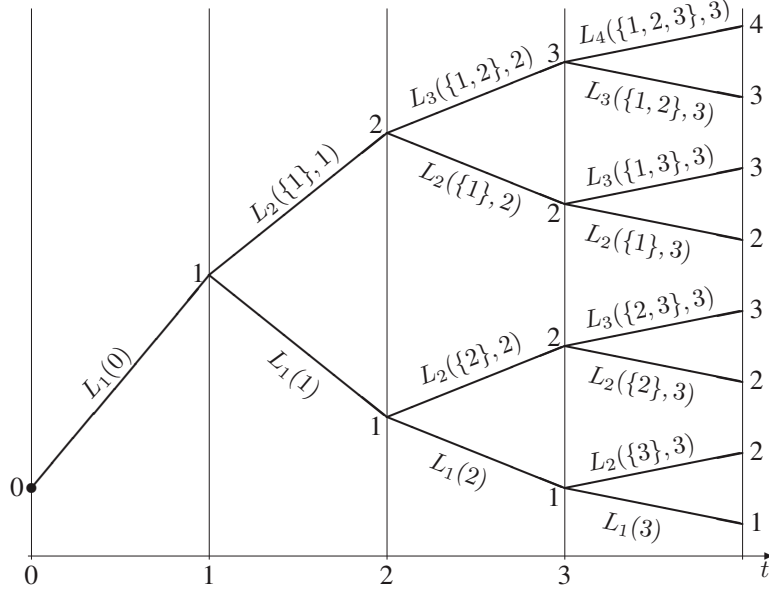


Figure 1: The events tree

Caption:

$L_m(\tau_m, t)$: per stage
payoff
 m : number of players
 τ_m : sequence of past
arrival times
 t : current time

3.1 Finite horizon

To better take into account the different arrival times of the players, we introduce an explicit discount factor $r \leq 1$, and write the payoff of the n -th player arrived as

$$\Pi_n(\tau_n) = \sum_{t=t_n}^T r^{t-t_n} L_{m(t)}(\tau_{m(t)}, t).$$

Both $m(t)$ and τ_m are random variables. Hence we let

$$\Pi_n^e(\tau_n) = \mathbb{E}\Pi_n.$$

We will prove the theorem:

Theorem 1 *The expected payoff of the discrete time, finite horizon game for player 1 is*

$$\Pi_1^e = \sum_{t=t_1}^T [r(1-p)]^{t-t_1} \sum_{m=1}^{t-t_1+1} \left(\frac{p}{1-p} \right)^{m-1} M_m(t). \quad (1)$$

We observe that thus the dependence of the per stage payoff on the sequence of arrival times has been lumped into the sequence of M_m . This is directly related to the fact that, for a Bernoulli process, at the number m of positive events given, all sequences of m arrival times share the same (conditional) probability. Accordingly, the dependence of the expected payoff on the per stage probability of arrival p is through the probabilities $p^{m-1}(1-p)^{t-t_1-m+1}$ of having $m-1$ arrivals between t_1+1 and t .

The fact that Π_1^e is decreasing in p , necessarily true, is not obvious in the formula (it depends on the evaluation of the M_m to be seen further down). Yet, it appears clearly that the limit expected payoff as p goes to zero is just

$$\Pi_1^e = \sum_{t=t_1}^T r^{t-t_1} L_1(t),$$

as the term $m=1$ only remains in the sum over m .

Proof of the theorem The proof of the theorem is given in appendix A.1. It is based upon a rather classical backward induction argument, and interchanges of orders of summations in the resulting formulas.

We may also write the equivalent formula for the payoff of the m -th arriving player. We need to introduce extra notation. For $1 \leq m < k$:

$$\tau_m^k = (\tau_{m+1}, \dots, \tau_k), \quad \mathcal{T}_m^k(t_m, t) = \{\tau_m^k \mid t_m < t_{m+1} < \dots < t_k \leq t\}.$$

and

$$M_m^m(\tau_m, t) = L_m(\tau_m, t) \quad \text{and} \quad M_m^k(\tau_m, t) = \sum_{\tau_m^k \in \mathcal{T}_m^k(t)} L_k(\tau_m, \tau_m^k, t).$$

Given m and t_m , the maximum possible number of players is $m + T - t_m$. By a computation similar to the one leading to the theorem, we get:

Corollary 1 *The payoff of the m -th arriving player, given the sequence of arrival times τ_m , is*

$$\Pi_m^e(\tau_m) = \sum_{\ell=0}^{T-t_m} \left(\frac{p}{1-p} \right)^\ell \sum_{t=t_m+\ell}^T [(1-p)r]^{t-t_m} M_m^{m+\ell}(\tau_m, t),$$

or, equivalently

$$\Pi_m^e(\tau_m) = \sum_{t=t_m}^T [(1-p)r]^{t-t_m} \sum_{\ell=0}^{t-t_m} \left(\frac{p}{1-p} \right)^\ell M_m^{m+\ell}(\tau_m, t). \quad (2)$$

3.2 Infinite horizon

The same set up is used to consider the game with an infinite horizon, *i.e.* a payoff

$$\Pi_1 = \sum_{t=0}^{\infty} r^t L_{m(t)}(\tau_{m(t)}, t).$$

We need two new definitions:

Definition 1 *The sequence of functions $\{L_m\}$ is said to be*

1. *uniformly bounded if it has a uniform bound, denoted L :*

$$\exists L > 0 : \forall t \in \mathbb{R}_+, \forall m \in \mathbb{N}, \forall \tau_m \in \mathcal{T}_m(t), \quad |L_m(\tau_m, t)| \leq L,$$

2. *exponentially bounded if each L_m is uniformly bounded, but that bound is allowed to vary exponentially with m :*

$$\exists L > 0 : \forall t \in \mathbb{R}_+, \forall m \in \mathbb{N}, \forall \tau_m \in \mathcal{T}_m(t), \quad |L_m(\tau_m, t)| \leq L^m.$$

The relationship of these properties depends on the position of L with respect to 1:

Remark 1

1. *If the sequence $\{L_m\}$ is uniformly bounded, it is also exponentially bounded, thus uniformly bounded is a stronger property than exponentially bounded. Indeed, let L be the uniform bound, the sequence is exponentially bounded by L^m if $L > 1$, by $1^m = 1$ if $L \leq 1$.*
2. *If the sequence $\{L_m\}$ is exponentially bounded by L^m with $L < 1$, it is also uniformly bounded by L . However, if $L > 1$, it may not be uniformly bounded.*

We prove the following theorem:

Theorem 2 *If the sequence $\{L_m\}$ is exponentially bounded by L^m and $r < 1/L$, or if it is uniformly bounded, then the expected payoff of the infinite horizon game is given by*

$$\Pi_1^e = \sum_{t=0}^{\infty} ((1-p)r)^t \sum_{m=1}^{t+1} \left(\frac{p}{1-p} \right)^{m-1} M_m(t). \quad (3)$$

Proof The proof, given in appendix A.2, is via taking the limit in formula 1, and carefully checking that the infinite series converge.

We derive the following consequences:

Corollary 2

- If the sequence $\{L_m\}$ is exponentially bounded by L^m with $1 \leq L < 1/r$, then

$$|\Pi_1^e| \leq \frac{L}{1 - rL},$$

- if the sequence $\{L_m\}$ is uniformly bounded by L , then

$$|\Pi_1^e| \leq \frac{L}{1 - r},$$

- if the L_m are constant and equal to L , and $p = 0$, then

$$\Pi_1^e = \frac{L}{1 - r}.$$

We may similarly extend formula (2) to an infinite series, which converges under the same conditions:

Corollary 3 *The infinite horizon payoff of the m -th arriving player is, under the same conditions as in theorem 2:*

$$\Pi_m^e(\tau_m) = r^{t_m} \sum_{t=t_m}^{\infty} ((1-p)r)^{t-t_m} \sum_{\ell=0}^{t-t_m} \left(\frac{p}{1-p} \right)^\ell M_m^{m+\ell}(\tau_m, t).$$

4 Applications to oligopoly theory

We consider identical firms entering a market. The game will be played over an infinite horizon, with a discount factor r . We investigate four different market structures, differing in the type of equilibrium —cartel or competitive à la Cournot— hypothesized at each stage, and in the behavior of the incumbents and of the new entrant each time one arrives. The four possible market structures will be called “cartel-cartel”, “cartel-Stackelberg”, “Cournot-Cournot”, and “Cournot-Stackelberg”, and will be explained in more detail below.

The per stage profits for the various equilibria considered are computed in the appendix B.

4.1 Cartel-cartel, or equally sharing a fixed revenue flux

In this simple case, the incumbents form a cartel, and each arriving new firm enters the cartel. If the optimum per stage profit feasible by a lone player on this market is c , we have $L_m(\tau_m, t) = c/m$. We assert

Theorem 3 *In the cartel-cartel market structure, the expected payoff to the first player is*

$$\Pi_1^e = \frac{c}{pr} \ln \left(1 + \frac{pr}{1-r} \right). \quad (4)$$

We are therefore able to get a closed form formula. It exhibits the limit expected payoff $\Pi_1^e = c/(1-r)$ as p goes to zero, and $\Pi_1^e \rightarrow \infty$ as $r \rightarrow 1$ for fixed p . That it is decreasing with p can be seen in the fact that the function $x \mapsto \ln(1+x)/x$ is decreasing for $x \in [0, 1]$.

Proof of theorem 3 The cardinal of the set $\mathcal{T}_m(t)$ is

$$|\mathcal{T}_m(t)| = \binom{t}{m-1} = \frac{t!}{(m-1)!(t-m+1)!}$$

we obtain

$$M_m(t) = c \frac{t!}{m!(t-m+1)!} = \frac{c}{t+1} \binom{t+1}{m},$$

and hence, applying formula (3),

$$\begin{aligned} \Pi_1^e &= c \sum_{t=0}^{\infty} r^t (1-p)^t \frac{1}{t+1} \sum_{m=1}^{t+1} \binom{t+1}{m} \left(\frac{p}{1-p} \right)^{m-1} \\ &= \frac{c}{pr} \sum_{t=0}^{\infty} \frac{r^{t+1} (1-p)^{t+1}}{t+1} \left[\left(1 + \frac{p}{1-p} \right)^{t+1} - 1 \right]. \end{aligned}$$

It suffices to recognize the identity, for $x < 1$:

$$\sum_{n=1}^{\infty} \frac{x^n}{n} = -\ln(1-x)$$

to conclude the proof. ■

It is possible to compute the expected payoff of the m -th arriving player, but we see little hope of simplifying it beyond the sheer repetition of the formula:

$$\Pi_m^e(\tau_m) = cr^{t_m-t_1} \sum_{t=t_m}^{\infty} [r(1-p)]^{t-t_m} \sum_{\ell=0}^{t-t_m} \frac{1}{m+\ell} \frac{(t-t_m)!}{\ell!(t-t_m-\ell)!} \left(\frac{p}{1-p} \right)^\ell,$$

or the slightly more appealing formula, that exhibits the independence of Π_m^e from τ_m :

Corollary 4 *The payoff of the m -th arriving player in the cartel-cartel market structure is*

$$\Pi_m^e(\tau_m) = cr^{t_m-t_1} \sum_{n=0}^{\infty} r^n n! \sum_{\ell=0}^n \frac{1}{m+\ell} \frac{p^\ell}{\ell!} \frac{(1-p)^{n-\ell}}{(n-\ell)!}.$$

4.2 Cartel-Stackelberg

In this market structure, in the absence of a new entrant, firms form a cartel. However, whenever a new firm enters the market, the incumbents still play as a cartel, but act as leaders imposing their strategy on the new entrant who acts as a follower, in a Stackelberg scheme. The table in the appendix shows that we may equivalently assume that the incumbents do not take immediate notice of the new entrant, or do not take it seriously, and continue with the same production level as before.

After a first step in that configuration, the new entrant joins the cartel for the rest of the game. It is not its best interest. But it may be coerced to do so by the other players who threat to revert otherwise to an all Cournot-Nash equilibrium. As soon as $m > 6$, they incur a loss in doing so, but not as much as the new entrant. It is therefore a credible threat.

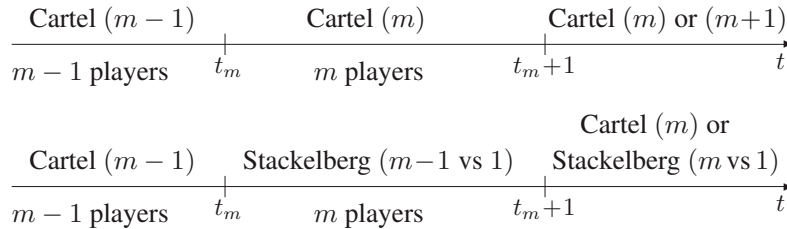


Figure 2: The time diagrams of the Cartel-Cartel (top) and Cartel-Stackelberg (bottom) market structures

The difference with the cartel-cartel time structure is displayed by the simple time diagrams of Figure 2, where we have assumed, for clarity, that $t_{m-1} < t_m - 1$, i.e. there was no new entrant at time $t_m - 1$.

According to the table in the appendix, we have, for some positive number c :

$$L_m(\tau_m, t_m) = \frac{c}{2(m-1)},$$

$$\forall t \in [t_m + 1, t_{m+1} - 1], L_m(\tau_m, t) = \frac{c}{m}.$$

We state the following theorem:

Theorem 4 *The expected payoff of the cartel-Stackelberg scheme for the first player is*

$$\Pi_1^e = c \left[1 + \frac{2-p}{2p} \ln \left(1 + \frac{pr}{1-r} \right) \right].$$

Although slightly more complicated than in the cartel-cartel scheme, this formula shares with the former one most of its characters: a closed form formula, exhibiting even more clearly its decreasing character with p , and converging to $V = c/(1-r)$ as $p \rightarrow 0$, and $V \rightarrow \infty$ as $r \rightarrow 1$.

Proof The proof, given in appendix A.4 involves a careful analysis of the combinatorics of the problem, and application of the previous formulas.

Finally, the payoff to the m -th arriving player at time t_m can be derived from the corresponding formula of the simple sharing problem, just correcting for the fact that at its first step, it earns $c/4$ rather than c/m .

Corollary 5 *The payoff of the m -th arriving player in the cartel-Stackelberg market structure is*

$$\Pi_m^e(t_m) = r^{t_m - t_1} \left[\frac{c}{4} - \frac{c}{m} + c \sum_{n=0}^{\infty} r^n n! \sum_{\ell=0}^n \frac{1}{m+\ell} \frac{p^\ell}{\ell!} \frac{(1-p)^{n-\ell}}{(n-\ell)!} \right].$$

4.3 Cournot-Cournot

We now assume that at each time step, the firms compete in a Cournot fashion. With reference to the diagrams of Figure 2, we now have that of Figure 3

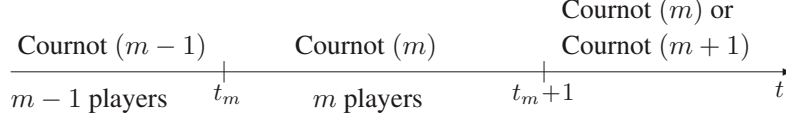


Figure 3: The time diagram of the Cournot-Cournot market structure

As a consequence, there is a positive number C such that

$$L_m = \frac{C}{(m+1)^2}.$$

We also assume a discount factor r as in the above theory. A direct application of the general theory leads to

Theorem 5 *The expected payoff of the first player in the Cournot-Cournot market structure is*

$$\Pi_1^e = C \sum_{t=0}^{\infty} r^t (1-p)^t \sum_{n=0}^t \left(\frac{p}{1-p} \right)^n \frac{t!}{(t-n)!n!} \frac{1}{(n+2)^2}. \quad (5)$$

This is not a very appealing formula. Yet, it is easy to program, even on a spreadsheet, to get a numerical approximation, computing the terms recursively. Write

$$u(t) = r^t (1-p)^t,$$

$$v_n(t) = \left(\frac{p}{1-p} \right)^n \frac{t!}{(t-n)!n!} \frac{1}{(n+2)^2},$$

$$\Pi_{1,C}^e = C \sum_{t=0}^{\infty} u(t) \sum_{n=0}^t v_n(t),$$

and compute the u and v_n according to the recursions $u(0) = 1$, $v_0(t) = 1/4$, $\forall n > t$, $v_n(t) = 0$, and, for $n \leq t$, $t \geq 1$:

$$u(t) = r(1-p)u(t-1),$$

$$v_n(t) = \frac{p}{1-p} \left(\frac{n+1}{n+2} \right)^2 v_{n-1}(t-1) + v_n(t-1).$$

Computing up to the fiftieth term, with $r = .8$, we found, for $p = 1/4$, $\Pi_{1,C}^e = .822 C$, for $p = 1/2$, $\Pi_{1,C}^e = .625 C$, and for $p = 3/4$, $\Pi_{1,C}^e = .508 C$. (To get the same precision, the computation requires the more terms that p is smaller. With 50 terms, for $p = 0$, we get $\Pi_{1,C}^e = 1.249986$ instead of the theoretical 1.25. These computations were performed on a spreadsheet LibreOffice.)

4.4 Cournot-Stackelberg scheme

We consider a scheme similar to that of the cartel-Stackelberg, but where firms compete à la Cournot-Nash instead of forming a cartel, both in the absence of a newcomer, and within the group of incumbents when a new firm enters the market. As in the cartel-Stackelberg structure, the behavior of the incumbents may be explained as ignoring the newcomer at the first step. Then (as soon as $m \geq 2$), all players profit from a reversal to the all Cournot-Nash equilibrium.

The time diagram is now as in Figure 4

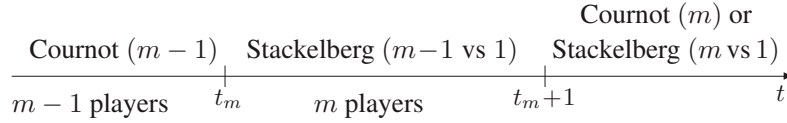


Figure 4: The time diagram of the Cournot-Stackelberg market structure

We therefore have

$$L_m(\tau_m, t_m) = \frac{C}{2m^2},$$

$$\forall t \in [t_m + 1, t_{m+1} - 1], L_m(\tau_m, t) = \frac{C}{(m+1)^2}.$$

The analysis is completely similar to that of the cartel-Stackelberg scheme, replacing the L_m as necessary, and leads to the following theorem:

Theorem 6 *Let $\Pi_{1,C}^e$ be the Cournot expected payoff (5). The expected payoff of the Cournot-Stackelberg scheme for the first player is*

$$\Pi_{1,CS}^e = \frac{r(2-p)}{2} \Pi_{1,C}^e + C.$$

We may notice that the relationship of this formula to the preceding one is the same as the corresponding one in the cartel case. It preserves the same properties as outlined above.

5 Conclusion

In this paper, we have analyzed a model where there is, for everyone, an unknown number of players. We have assumed that all the players are *clones* (since they

adopt the same strategy and had the same payoff at each period of time) and that demand market and entry device (Bernoulli process) are common knowledge. Players' payoff differ from each other only due to the time during which they are in the market. In this setting we are able to calculate the payoff of each player in finite and infinite horizon games. Four simple examples in microeconomics show that the theory is indeed operative. The formulas obtained are not very simple, but they allow for efficient numerical implementations.

Yet, we may point out a set of weaknesses of the theory, each of which points to possible further investigations.

- We are limited to a *clone economy*. Dealing with an unlimited number of players, the game has to be anonymous. But we might want to have several *classes* of players, such as done in many studies of population uncertainty and anonymous games.
- To keep with the simple formulas of the Bernoulli process, we impose a constant entry probability p . It would be more realistic to have it depend on the number of players already on the market, as the benefit of entry decreases with that number. This is chiefly true in our examples where the size of the market is constant. This would lead to exceedingly complex formulas, hardly implementable, except may be if we let p be a constant up to a certain number of players and drop to zero afterward.
- In our four examples, the players' per stage payoff depend on the number of players, and only very mildly on the sequence of past arrival times. (In the one-step Stackelberg games, where we need to distinguish the case $t = t_m$ from the case $t > t_m$.) A more complex dependence such as alluded to at the beginning of section 3, say because a resource is consumed by the players, requires, to remain manageable by our theory, that it be explicit enough to let us compute the sums M_m . A rather severe restriction on the models we are able to deal with.
- We do not allow for players randomly leaving the market. Yet, in many applications, this would be more realistic.

But as it is, the theory can probably be used as a Rubinstein “fable” ([Rubinstein, 2006, Rubinstein, 2012]) to investigate some real life economic problems.

References

- [Breton and Keoula, 2012] Breton, M. and Keoula, M. Y. (2012). Farsightedness in a coalitional great fish war. *Environmental and Resource Economics*, 51:297–315.
- [Breton et al., 2010] Breton, M., Sbragia, L., and Zaccour, G. (2010). A dynamic model for international environmental agreements. *Environmental and Resource Economics*, 45:25–48.
- [Davis, 1985] Davis, M. H. A. (1985). Control of piecewise-deterministic processes via discrete-time dynamic programming. In *Stochastic Differential Systems*, volume 78 of *Lecture Notes in Control and Information Sciences*, pages 140–150. Springer.
- [De Sinopoli et al., 2014] De Sinopoli, F., Meroni, C., and Pimienta, C. (2014). Strategic stability in Poisson games. *Journal of Economic Theory*, 153:46–63.
- [De Sinopoli and Pimienta, 2009] De Sinopoli, F. and Pimienta, C. (2009). Undominated (and) perfect equilibria in Poisson games. *Games and Economic Behavior*, 66:775–784.
- [Haurie et al., 1994] Haurie, A., Leizarowitz, A., and van Delft, C. (1994). Boundedly optimal control of piecewise deterministic systems. *European journal of Operational Research*, 73:237–251.
- [Levin and Ozdenoren, 2004] Levin, D. and Ozdenoren, E. (2004). Auctions with uncertain number of bidders. *Journal of Economic Theory*, 118:229–251.
- [Makris, 2008] Makris, M. (2008). Complementarities and macroeconomics: Poisson games. *Games and Economic Behavior*, 62:180–189.
- [Matthews, 1987] Matthews, S. (1987). Comparing auctions for risk averse buyers: a buyer’s point of view. *Econometrica*, 55:633–646.
- [McAfee and McMillan, 1987a] McAfee, R. P. and McMillan, J. (1987a). Auctions with a stochastic number of bidders. *Journal of Economic Theory*, 43:1–19.
- [McAfee and McMillan, 1987b] McAfee, R. P. and McMillan, J. (1987b). Auctions with entry. *Economic Letters*, 23:343–347.
- [Milchtaich, 2004] Milchtaich, I. (2004). Random-player games. *Games and Economic Behavior*, 47:353–388.
- [Münster, 2006] Münster, J. (2006). Contests with an unknown number of contestants. *Public Choice*, 129:353–368.
- [Myerson, 1998a] Myerson, R. B. (1998a). Extended Poisson games and the Condorcet jury theorem. *Games and Economic Behavior*, 25:111–131.
- [Myerson, 1998b] Myerson, R. B. (1998b). Population uncertainty and Poisson games. *International Journal of Game Theory*, 27:375–392.
- [Myerson, 2000] Myerson, R. B. (2000). Large Poisson games. *Journal of Eco-*

- nomic Theory*, 94:7–45.
- [Neyman and Sorin, 2003] Neyman, A. and Sorin, S. (2003). *Stochastic games*. NATO ASI series. Kluwer academic publishers.
- [Nowak et al., 2013] Nowak, A., Sloan, E., and Sorin, S. (2013). Special issue on stochastic games. *Dynamic Games and Applications*, 3.
- [Östling et al., 2011] Östling, R., Tao-Yi Wang, J., Chou, E. Y., and Camerer, C. F. (2011). Testing game theory in the field: Swedish lupi lottery games. *American Economic Journal: Microeconomics*, 3:1–33.
- [Piccione and Tan, 1996] Piccione, M. and Tan, G. (1996). A simple model of expert and non expert bidding in first price auctions. *Journal of Economic Theory*, 70:501–515.
- [Ritzberger, 2009] Ritzberger, K. (2009). Price competition with population uncertainty. *Mathematical Social Sciences*, 58:145–157.
- [Rubinstein, 2006] Rubinstein, A. (2006). Dilemmas of an economic theorist. *Econometrica*, 74:865–883.
- [Rubinstein, 2012] Rubinstein, A. (2012). *Economic Fables*. Open Books Publishers.
- [Rubio and Ulph, 2007] Rubio, J. S. and Ulph, A. (2007). An infinite-horizon model of dynamic membership of international environmental agreements. *Journal of Environmental Economics and Management*, 54:296–310.
- [Salo and Weber, 2007] Salo, A. and Weber, M. (2007). Ambiguity aversion in first-price sealed bid auctions. *Journal of Risk and Uncertainty*, 11:123–137.
- [Samuelson, 1985] Samuelson, W. F. (1985). Competitive bidding with entry costs. *Economic Letters*, 17:53–57.
- [Sworder, 1969] Sworder, D. D. (1969). Feedback control of a class of linear systems with jump parameters. *IEEE Transactions on Automatic Control*, AC-14:9–14.

A Proofs of the theorems

A.1 Proof of theorem 1

We remark the maximum number of players is $T - t_1 + 1$, can only be attained at time T , and only if a player has arrived at each instant of time. We then have

$$\Pi_{T-t_1+1}^e(t_1, t_1 + 1, \dots, T) = L_{T-t_1+1}(T) = M_{T-t_1+1}(T).$$

For $m \leq T - t_1$ and a compatible τ_m , we have

$$J_m(\tau_m) = \begin{cases} \sum_{t=t_m}^T r^{t-t_m} L_m(\tau_m, t) & \text{if } t_{m+1} > T, \\ \sum_{t=t_m}^{t_{m+1}-1} r^{t-t_m} L_m(\tau_m, t) + r^{t_{m+1}-t_m} J_{m+1}(\tau_m, t_{m+1}) & \text{if } t_{m+1} \leq T. \end{cases}$$

Now, $t_{m+1} > T$ with a probability $(1-p)^{T-t_m}$, and the occurrence of a given $t_{m+1} \leq T$ has a probability $p(1-p)^{t_{m+1}-t_m-1}$. Hence, writing t_+ for t_{m+1} :

$$\begin{aligned} \Pi_m^e(\tau_m) &= (1-p)^{T-t_m} \sum_{t=t_m}^T r^{t-t_m} L_m(\tau_m, t) + \\ &\sum_{t_+=t_m+1}^T p(1-p)^{t_+-t_m-1} \left[\sum_{t=t_m}^{t_+-1} r^{t-t_m} L_m(\tau_m, t) + r^{t_+-t_m} \Pi_{m+1}^e(\tau_m, t_+) \right]. \end{aligned}$$

Introduce the notation $q = 1 - p$. Interchanging the order of summation,

$$\begin{aligned} \Pi_m^e(\tau_m) &= q^{T-t_m} \sum_{t=t_m}^T r^{t-t_m} L_m(\tau_m, t) + \\ &\sum_{t=t_m}^{T-1} r^{t-t_m} L_m(\tau_m, t) \sum_{t_+=t+1}^T p q^{t_+-t_m-1} + \sum_{t_+=t_m+1}^T p q^{t_+-t_m-1} r^{t_+-t_m} \Pi_{m+1}^e(\tau_m, t_+). \end{aligned}$$

Using classical formulas for the sum of a geometric series, and regrouping terms we obtain for $m \leq T - t_1$:

$$\Pi_m^e(\tau_m) = \sum_{t=t_m}^T q^{t-t_m} \left[r^{t-t_m} L_m(\tau_m, t) + p r^{t_{m+1}-t_m} \Pi_{m+1}^e(\tau_m, t+1) \right],$$

being agreed that $\Pi_{m+1}^e(\tau_m, T+1) = 0$. A more useful form of this formula for the sequel is as follows:

$$\Pi_m^e(\tau_m) = \sum_{t=t_m}^T (qr)^{t-t_m} L_m(\tau_m, t) + \frac{p}{q} \sum_{t_{m+1}=t_m+1}^T (qr)^{t_{m+1}-t_m} \Pi_{m+1}^e(\tau_m, t_{m+1}), \quad (6)$$

where the second term of the right hand side is absent if $t_m = T$.

Remark 2 It is a not-so-easy calculation to check that, if $L_m(\tau_m, t) \leq L$ and $\Pi_{m+1}^e(\tau_m, t_{m+1}) \leq (T - t_{m+1} + 1)L$, then this formula implies, as it should, $\Pi_m^e(\tau_m) \leq (T - t_m + 1)L$.

A hint about how to make the above check is as follows: in the second term of the formula, write

$$\sum_{t=t_m+1}^T q^{t-t_m}(T-t+1) = \left(T - t_m + 1 - (1-p) \frac{d}{dp} \right) \sum_{t=t_m+1}^T (1-p)^{t-t_m},$$

and use the classic formula for the sum of the (finite) geometric series.

We use formula (6) recursively: write first Π_1^e as a function of Π_2^e , and using again the same formula substitute for Π_2^e as a function of Π_3^e , and again for Π_3^e as a function of Π_4^e . Then interchange the order of the summations, placing them in the order t, t_2, t_3, t_4 . In the following formula, every time the lower bound of a sum is larger than the upper bound, the term is just absent. We obtain

$$\begin{aligned} \Pi_1^e &= \sum_{t=t_1}^T (qr)^{t-t_1} L_1(t_1, t) + \frac{p}{q} \sum_{t=t_1+1}^T (qr)^{t-t_1} \sum_{t_2=t_1+1}^t L_2(t_1, t_2, t) \\ &+ \left(\frac{p}{q} \right)^2 \sum_{t=t_1+2}^T (qr)^{t-t_1} \sum_{t_2=t_1+1}^T \sum_{t_3=t_2+1}^t L_3(t_1, t_2, t_3, t) \\ &+ \left(\frac{p}{q} \right)^3 \sum_{t_2=t_1+1}^T \sum_{t_3=t_2+1}^T \sum_{t_4=t_3+1}^T (qr)^{t_4-t_1} \Pi_4^e(t_1, t_2, t_3, t_4). \end{aligned}$$

Continuing in the same way up to $m = T - t_1 + 1$, we obtain

$$\Pi_1^e = \sum_{m=1}^{T-t_1} \left(\frac{p}{q} \right)^{m-1} \sum_{t=t_1+m-1}^T (qr)^{t-t_1} M_m(t) + \left(\frac{p}{q} \right)^{T-t_1} (qr)^{T-t_1} L_{T-t_1+1}(T).$$

The last term can be identified as the term $m = T - t_1 + 1$ of the first sum, as the range of t in the embedded (second) sum is limited to $t = T$, and we have seen that $L_{T-t_1+1}(T) = M_{T-t_1+1}(T)$. It suffices now to shift m by one unit to obtain the formula

$$\Pi_1^e = \sum_{m=1}^{T-t_1+1} \left(\frac{p}{1-p} \right)^{m-1} \sum_{t=t_1+m-1}^T [(1-p)r]^{t-t_1} M_m(t).$$

And interchanging a last time the order of the summations, formula (1). ■

Remark 3 As a consequence of formula (1), if for some fixed L , for all m, τ_m and t , $L_m(\tau_m, t) = L$, then $\Pi_1^e = [(1 - r^{T-t_1+1})/(1 - r)]L$ (whose limit as $r \rightarrow 1$ is $(T - t_1 + 1)L$), and if $L_m(\tau_m, t) \leq L$, then Π_1^e is bounded above by that number.

A.2 Proof of theorem 2

We start with formula (1) where we set $t_1 = 0$, and recall by a superscript (T) that it is a formula for a finite horizon T , the horizon we want to let go to infinity:

$$\Pi_1^{e(T)} = \sum_{t=0}^T (1-p)^t r^t \sum_{m=1}^{t+1} \left(\frac{p}{1-p} \right)^{m-1} M_m(t).$$

The only task left to prove the theorem is to show that the series obtained as $T \rightarrow \infty$ converges absolutely. To do this, we need an evaluation of $M_m(t)$. Observe that the cardinal of the set $\mathcal{T}_m(t)$ is simply the combinatorial coefficient

$$|\mathcal{T}_m(t)| = \binom{t}{m-1} = \frac{t!}{(m-1)!(t-m+1)!}.$$

As a consequence, if $|L_m| \leq L^m$, we have

$$|M_m(t)| \leq \binom{t}{m-1} L^m$$

and

$$\left| \sum_{m=1}^{t+1} \left(\frac{p}{1-p} \right)^{m-1} M_m(t) \right| \leq \sum_{m=1}^{t+1} \left(\frac{p}{1-p} \right)^{m-1} |M_m(t)| \quad (7)$$

$$\leq \sum_{m=1}^{t+1} \binom{t}{m-1} \left(\frac{p}{1-p} \right)^{m-1} L^m \quad (8)$$

$$= L \left(\frac{p}{1-p} L + 1 \right)^t. \quad (9)$$

Therefore,

$$(1-p)^t r^t \sum_{m=1}^{t+1} \left(\frac{p}{1-p} \right)^{m-1} |M_m(t)| \leq L r^t (pL + 1 - p)^t.$$

The series converges absolutely provided that

$$r(p(L-1) + 1) < 1,$$

which is always true if $L \leq 1$, and ensured for all $p \leq 1$ if $rL < 1$. This proves the theorem for the case exponentially bounded.

In the case uniformly bounded, with $|L_m| \leq L$, we obtain similarly

$$|M_m(t)| \leq \binom{t}{m-1} L$$

and

$$(1-p)^t r^t \sum_{m=1}^{t+1} \left(\frac{p}{1-p} \right)^{m-1} |M_m(t)| \leq L r^t (1-p)^t \left(\frac{p}{1-p} + 1 \right)^t = L r^t,$$

and the series is always absolutely convergent. \blacksquare

A.4 Proof of theorem 4

We aim to apply formula (3). The term $t = 0$ requires a special treatment: the only term in the sum over m is $m = 1$ and $M_1(0) = L_1(0) = c$. For $t > 0$, we have three cases:

1. For $m = 1$, there has not been any new entrant, therefore $L_1(t) = c$.
2. For $1 < m < t + 1$, we sum first over the τ_m such that $t_m < t$, i.e. $\tau_m \in \mathcal{T}_m(t-1)$, then over the τ_m such that $t_m = t$; there the sum is over the values of $\tau_{m-1} \in \mathcal{T}_{m-1}(t-1)$.
3. For $m = t + 1$, there have been new entrants at each time step. Therefore $L_{t+1}(\tau_{t+1}, t) = c/2t$. We summarize this in the following calculation:

$$\begin{aligned} \text{for } m = 1, \quad M_1(t) &= c = c \frac{(t-1)!}{(m-1)!(t-m+1)!} \frac{1}{m}, \\ \text{for } 1 < m < t + 1, \quad M_m(t) &= \sum_{\tau_m \in \mathcal{T}_m(t-1)} \frac{c}{m} + \sum_{\tau_{m-1} \in \mathcal{T}_{m-1}(t-1)} \frac{c}{2(m-1)} \\ &= c \frac{(t-1)!}{(m-1)!(t-m+1)!} \frac{1}{m} \\ &\quad + c \frac{(t-1)!}{(m-2)!(t-m+1)!} \frac{1}{2(m-1)}, \\ \text{for } m = t + 1, \quad M_{t+1}(t) &= \frac{c}{2t} = c \frac{(t-1)!}{(m-2)!(t-m+1)!} \frac{1}{2(m-1)}. \end{aligned}$$

We therefore obtain, for $t > 0$:

$$\begin{aligned} \sum_{m=1}^{t+1} \left(\frac{p}{q}\right)^{m-1} M_m(t) &= \sum_{m=1}^t \left(\frac{p}{q}\right)^{m-1} \frac{(t-1)!}{(m-1)!(t-m)!} \frac{c}{m} \\ &\quad + \frac{1}{2} \sum_{m=2}^{t+1} \left(\frac{p}{q}\right)^{m-1} \frac{(t-1)!}{(m-2)!(t-m+1)!} \frac{c}{m-1} \\ &= \left(1 + \frac{p}{2q}\right) \sum_{m=1}^t \left(\frac{p}{q}\right)^{m-1} \frac{(t-1)!}{(m-1)!(t-m)!} \frac{c}{m}. \end{aligned}$$

Finally, summing over t as in formula (3), without forgetting the term $t = 0$,

$$\Pi_1^e = c + \left(1 + \frac{p}{2q}\right) \sum_{t=1}^{\infty} r^t q^t \sum_{m=1}^t \left(\frac{p}{q}\right)^{m-1} \frac{(t-1)!}{(m-1)!(t-m)!} \frac{c}{m}.$$

It suffices to take out one power of rq from the sum over t , shift the summation index by one unit, recognize the expected payoff of the simple sharing scheme and replace it by formula (4) to prove the theorem.

B Static equilibria

We consider n identical firms (a clone economy) sharing a market with a linear inverse demand function, linking the price P to the total production level Q as

$$P = b - aQ.$$

Production costs have been lumped into b and so doing normalized at zero. We compute the optimal production level Q , and resulting price P and profit Π of each firm for various equilibria.

B.1 Cartel

In a pure cartel, firms behave as a single player, only sharing the optimal production level equally among them. Let Q be that level. The profit of the (fictitious) single player is

$$\Pi = Q(b - aQ) = -a \left(Q - \frac{b}{2a}\right)^2 + \frac{b^2}{4a}.$$

Hence the optimal production level is $Q = b/(2a)$, to be equally divided among the firms, as well as the profit $\Pi = b^2/(4a)$. The price is then $P = b/2$, and the

individual production level q and profit Π_i are

$$q = \frac{b}{2an}, \quad \Pi_i = \frac{b^2}{4an}.$$

B.2 Cartel-Stackelberg

We investigate the case where $n - 1$ firms form a cartel, behaving as a leader *vis à vis* one firm acting as a follower.

Let q_L be the quantity produced by each incumbent, q_F that of the follower. Hence $Q = (n - 1)q_L + q_F$. The follower's profit is

$$\Pi_F = q_F[b - a(n - 1)q_L - aq_F] = -a \left[q_F^2 - \left(\frac{b}{a} - (n - 1)q_L \right) q_F \right]$$

hence

$$\Pi_F = -a \left[q_F - \frac{1}{2} \left(\frac{b}{a} - (n - 1)q_L \right) \right]^2 + \frac{a}{4} \left(\frac{b}{a} - (n - 1)q_L \right)^2.$$

Therefore, the optimal reaction curve q_F as a function of q_L is

$$q_F = \frac{1}{2} \left(\frac{b}{a} - (n - 1)q_L \right).$$

With such a strategy, each incumbents' profit is

$$\begin{aligned} \Pi_L &= q_L \left[b - a(n - 1)q_L - \frac{1}{2}b - \frac{1}{2}a(n - 1)q_L \right] \\ &= -\frac{a(n - 1)}{2} \left[q_L^2 - \frac{b}{a(n - 1)}q_L \right] \\ &= -\frac{a(n - 1)}{2} \left[q_L - \frac{b}{2a(n - 1)} \right]^2 + \frac{b^2}{8a(n - 1)}. \end{aligned}$$

Therefore, the optimal production level of each incumbent and their profit are

$$q_L = \frac{b}{2a(n - 1)}, \quad \Pi_L = \frac{b^2}{8a(n - 1)}.$$

Placing this back into the optimal follower's reaction curve, its production level and profit are

$$q_F = \frac{b}{4a}, \quad \Pi_F = \frac{b^2}{16a}.$$

and the price of the commodity is $P = b/4$. All these results will be summarized in a table in the last section.

B.3 Cournot-Nash

We have n identical firms competing à la Cournot. The Cournot-Nash equilibrium is obtained as follows. Let q be the individual production level, therefore $Q = nq$, and P the resulting price:

$$P = b - naq.$$

The individual profit of player i is

$$\begin{aligned}\Pi_i &= q_i[b - a(q_i + (n-1)q)] \\ &= -a \left[q_i - \frac{1}{2} \left(\frac{b}{a} - (n-1)q \right) \right]^2 + \frac{a}{4} \left(\frac{b}{a} - (n-1)q \right)^2.\end{aligned}$$

It follows that the optimum q_i is

$$q_i = \frac{1}{2} \left(\frac{b}{a} - (n-1)q \right),$$

but we seek a symmetric equilibrium where $q_i = q$, and therefore

$$q = \frac{b}{(n+1)a}.$$

Placing this back into the law for P , we find

$$P = \frac{b}{(n+1)}, \quad \Pi = \frac{b^2}{a(n+1)^2}.$$

B.4 Cournot-Stackelberg

We finally consider $n-1$ firms competing à la Cournot-Nash within their group, producing a quantity q_L each, but that group behaving as a leader *vis à vis* a single follower, producing a quantity q_F . We therefore have

$$P = b - a(n-1)q_L - aq_F.$$

The calculations are similar to the previous ones. The follower's profit is therefore

$$\begin{aligned}\Pi_F &= -a \left[q_F^2 - \left(\frac{b}{a} - (n-1)q_L \right) q_F \right] \\ &= -a \left[q_F - \frac{1}{2} \left(\frac{b}{a} - (n-1)q_L \right) \right]^2 + \frac{a}{4} \left(\frac{b}{a} - (n-1)q_L \right)^2.\end{aligned}$$

Hence

$$q_F = \frac{1}{2} \left(\frac{b}{a} - (n-1)q_L \right), \quad \Pi_F = \frac{a}{4} \left(\frac{b}{a} - (n-1)q_L \right)^2.$$

With this strategy,

$$P = \frac{1}{2} (b - a(n-1)q_L) = \frac{1}{2} (b - a(n-2)q_L - aq_i)$$

Consequently, for player i , one of the leaders,

$$\begin{aligned} \Pi_i &= -\frac{a}{2} \left[q_i^2 - \left(\frac{b}{a} - (n-2)q_L \right) q_i \right] \\ &= -\frac{a}{2} \left[q_i - \frac{1}{2} \left(\frac{b}{a} - (n-2)q_L \right) \right]^2 + \frac{a}{8} \left(\frac{b}{a} - (n-2)q_L \right)^2. \end{aligned}$$

It follows that

$$\begin{aligned} q_i = q_L = \frac{1}{2} \left(\frac{b}{a} - (n-2)q_L \right) &\Rightarrow q_L = \frac{b}{an}, \\ \Pi_i = \frac{a}{8} \left(\frac{b}{a} - (n-2)q_L \right)^2 &\Rightarrow \Pi_i = \frac{b^2}{2an^2}, \end{aligned}$$

while $P = b/(2n)$, and

$$q_F = \frac{b}{2an} \quad \text{and} \quad \Pi_F = \frac{b^2}{4an^2}.$$

B.5 Summary

We regroup the results of these calculations in the following table:

Market structure	# of firms	Quantity	Price	Profit
Cartel	n	$\frac{b}{2an}$	$\frac{b}{2}$	$\frac{b^2}{4an}$
Cartel-Stackelberg	n	$\frac{3b}{4a}$	$\frac{b}{4}$	
each leader	$n - 1$	$\frac{b}{2a(n - 1)}$		$\frac{b^2}{8a(n - 1)}$
follower	1	$\frac{b}{4a}$		$\frac{b^2}{16a}$
Oligopoly à la Cournot	n	$\frac{b}{a(n + 1)}$	$\frac{b}{(n + 1)}$	$\frac{b^2}{a(n + 1)^2}$
Cournot-Stackelberg	n	$\frac{b(n + 1/2)}{an}$	$\frac{b}{2n}$	
each leader	$n - 1$	$\frac{b}{an}$		$\frac{b^2}{2an^2}$
follower	1	$\frac{b}{2an}$		$\frac{b^2}{4an^2}$

C Complexity

In this appendix, we undertake to count the number of arithmetic operations involved in computing Π_1^e for a finite horizon by four different methods: (direct) path enumeration, backward dynamic programming, path enumeration and forward dynamic programming, and use of formula (1). To ease the calculations, we let $t_1 = 1$, so that T is the number of time steps. We shall refer to the tree of Figure 1.

If we assume no regularity in their definition, the data is made of the collection of all $L_m(t)$, $t = 1, \dots, T$, that is as many numbers as there are branches in the tree, i.e.

$$\sum_{t=1}^T 2^{t-1} = 2^T - 1$$

numbers. As all must be used, there is no way in which a general method could

involve less than that number of arithmetic operations. Therefore, we expect a complexity of the order of 2^T (of the order of 10^6 for $T = 20$, and 10^{15} for $T = 50$, a completely unrealistic case !), and the difference between methods can only be in the coefficient multiplying that number.

C.1 Path enumeration

The tree counts 2^{T-1} paths from the root to a leaf. Let $\nu \in [1, 2^{T-1}]$ number them. We denote by π_ν the path number ν , and let m_ν be the number of player present at the end of path π_ν . Each path has a probability of being followed

$$\mathbb{P}(\pi_\nu) = p^{m_\nu-1}(1-p)^{T-m_\nu}.$$

Let $L_\nu(t)$ denote the $L_n(\tau_n, t)$ on path π_ν . Each path involves a payoff

$$\Pi_1(\pi_\nu) = \sum_{t=1}^T L_\nu(t).$$

And we have

$$\Pi_1^e = \sum_{\nu=1}^{2^{T-1}} \mathbb{P}(\pi_\nu) J(\pi_\nu). \quad (10)$$

A direct method of evaluating Π^e is therefore as follows:

1. compute the $\mathbb{P}(\pi_\nu)$ for each ν . The computation of each involves $T - 2$ multiplications². Therefore that step involves $2^{T-1}(T - 2)$ arithmetic operations.
2. Compute the $\Pi(\pi_\nu)$. Each involves $T - 1$ additions, therefore this step involves $2^{T-1}(T - 1)$ arithmetic operations.
3. Compute Π_1^e according to formula (10), involving 2^{T-1} multiplications and as many additions (-1), that is 2^T operations.

Therefore the total number of arithmetic operations is

$$N = 2^{T-1}(T - 2 + T - 1 + 2) = (T - \frac{1}{2})2^T,$$

that is of the order of $T2^T$.

C.2 Dynamic programming (DP)

Denote the nodes of the tree by the sequence $\sigma(t)$ of t indices, 0 or 1, the 1 denoting the times when an arrival occurred, a branch sloping up in our figure. (All

²we count powers as a sequence of multiplications

sequences $\sigma(t)$ begin with a one.) The possible successors of a given $\sigma(t)$ are $(\sigma(t), 0)$ and $(\sigma(t), 1)$, that we denote as

$$\sigma(t+1) = (\sigma(t), i(t)), \quad i(t) \in \{0, 1\}.$$

Denote by $L(\sigma(t))$ the L_m of the branch reaching the node $\sigma(t)$.

C.2.1 Backward DP

Let $V(\sigma(t))$ be the expected future payoff when at node $\sigma(t)$. It obeys the dynamic programming equation

$$V(\sigma(t)) = p[L(\sigma(t), 1) + V(\sigma(t), 1)] + (1-p)[L(\sigma(t), 0) + V(\sigma(t), 0)],$$

and $\Pi_1^e = V(\text{root}) = V(1) + L_1(1)$.

There are thus four arithmetic operations to perform at each node of the tree (not counting the leaves), that is

$$N = 4 \sum_{t=1}^T 2^{t-1} = 4 \times (2^T - 1)$$

arithmetic operations, i.e. of the order of 4×2^T .

C.2.2 Path enumeration and forward DP

This is a variant of the path enumeration method (10) on two counts:

1. Compute once each probability $p^{m-1}(1-p)^{T-m}$ and store it. This costs $T(T-2)$ arithmetic operations.
2. Compute the $\Pi(\pi_\nu)$ according to the forward dynamic programming method

$$\Pi(\sigma(t-1), i(t-1)) = \Pi(\sigma(t-1)) + L(\sigma(t-1), i(t-1)).$$

This is one addition at each node of the tree, i.e. 2^T operations.

It remains to implement formula (10), using 2^T arithmetic operations, for a total of $2^{T+1} + T(T-2)$. This is of the order of 2×2^T .

C.3 Using the $M_m(t)$

We rewrite formula (1) as

$$\Pi_1^e = \sum_{t=1}^T \sum_{m=1}^t p^{m-1} (1-t)^{t-m} M_m(t). \quad (11)$$

C.3.1 Computing the $M_m(t)$

The first task is to compute the collection of $M_m(t)$. For each t , there are 2^{t-1} $L_m(t)$ to combine in t terms, that is $2^{t-1} - t$ additions. There is no computation for the steps 1 and 2. The total number of additions there is therefore

$$\sum_{t=3}^T (2^{t-1} - t) = 2^T - \frac{T(T+1)}{2} - 1.$$

C.3.2 Computing the $p^{m-1}(1-p)^{t-m}$

We set

$$u_m(t) = p^{m-1}(1-p)^{t-m}.$$

We compute them according to the following method:

$$\begin{aligned} u_1(1) &= 1, \\ \forall t \in [2, T], \quad u_1(t) &= (1-p)u_1(t-1), \\ \forall m \in [1, t], \quad u_{m+1}(t) &= pu_m(t-1). \end{aligned}$$

Counting the arithmetic operations, this leads to $T - 1$ multiplications to compute the $u_1(t)$, and to

$$\sum_{t=3}^T (t-1) = \frac{T(T-1)}{2} - 3$$

multiplications to compute the rest of the $u_m(t)$. That is for this step

$$\frac{T(T-1)}{2} + T - 4$$

arithmetic operations.

C.3.3 Applying formula (11)

Finally, there are $T(T+1)/2$ terms in formula (11), each involving a multiplication and an addition, i.e. $T(T+1)$ arithmetic operations (minus one addition). Summing all steps, this is

$$N = 2^T + T^2 + T - 6$$

i.e. of the order of 2^T arithmetic operations, half as many as in the best DP method.

C.4 Conclusion

Our theory gives the fastest algorithm for this general “unstructured” case, half as many algebraic operations as in the next best. But of course, its main advantage is elsewhere, in allowing one to take advantage of any regularity in the definitions of the L_m , and also in allowing for closed formulas for the infinite horizon case. Formula (4) is a typical example.

A general remark is that going from the “direct” method, in $T2^T$ arithmetic operations to one with a constant coefficient, in a typical computer science tradeoff, we trade computer time for storage space. However, if the L_m need to be stored as data (as opposed to being given by some explicit formula), then in all the faster methods, they are used only once so that their memory space can be re-used for intermediate results.

D Dynamic Nash equilibrium

We have stressed that, in our applications, we seek an equilibrium at each time step. But our process is then a sequence of static equilibria, not a dynamic equilibrium. As stated in the main body, our theory does not address the determination of a dynamic equilibrium. In the finite horizon case, it is possible to extend the theory to encompass a dynamic equilibrium, computed retrogressively *à la* Isaacs, using (a discrete time parallel of) the theory of piecewise deterministic systems. (See [Sworder, 1969, Davis, 1985, Haurie et al., 1994]). But we see little applicability of that extended theory. We give a short account of it here for the sake of completeness.

Let $U(t)$ be a set of possible decisions at time t for the (identical) players. We consider situations where the focal player, say player 1, uses the decisions $a(t) \in U(t)$, and all other players decisions $b(t) \in U(t)$. This suffices according to Isaacs’ tenet of transition (Bellman’s optimality principle), because all players present at a given time face the same game problem regarding the future, regardless of when they arrived. Let $\alpha(t)$ be the sequence of decisions $u(s)$ of player 1 from $s = t_1$ up to $s = t - 1$, and $\beta(t)$ the sequence of decisions of all other players, also up to time $t - 1$. (And for $t \leq t_2$, we set $b(t) = \text{nil.}$) Let $L_m(\tau_m, \alpha(t), a, \beta(t), b, t)$ be the per stage payoff of the players. (It could also depend on $\alpha(t)$ and $\beta(t)$.)

Instead of the expected payoff Π_m^e , introduce a family of Value functions V_m defined for $t \geq t_m$:

$$V_m(\tau_m, \alpha(t), \beta(t), t).$$

Nash equilibrium strategies $a^*(\cdot)$ must satisfy

$$V_m(\{\alpha(t), a^*(t)\}, \{\beta(t), a^*(t)\}, t) = \max_{a \in U(t)} \left[L_m(\tau_m, \alpha(t), a, \beta(t), a^*(t)) + pV_{m+1}(\{\tau_m, t\}, \{\alpha(t), a\}, \{\beta(t), a^*(t)\}, t+1) + (1-p)V_m(\tau_m, \{\alpha(t), a\}, \{\alpha(t), a^*(t), t+1\}) \right].$$

the maximum being reached at $a = a^*$. This recursion, initialized at

$$V_m(\cdot, \cdot, \cdot, T) = 0,$$

may theoretically allow one to compute the Value functions V_m and the equilibrium decisions.

Of course, this simplifies if the dependence of the L_m on τ_m , α and β is through a state equation such as hinted at at the beginning of our section 3. But even so, we see little applicability of that theory.

One exception might be in the very un-natural case of an infinite horizon problem with such a state living in a finite dimensional (or better: finite) space, where U is constant, and where the L_m are of the form

$$L_m(\alpha(t), a, \beta(t), b, t) = r^t s^m \mathcal{L}(x(t), a_1, a_2, \dots, a_m).$$

Let $V_m(x(t), t) = r^t s^m W_m(x)$. Then the above equation takes the form of a modified Shapley equation:

$$W_m(x) = \max_{a \in U} \left[\mathcal{L}_m(x, a, a^*, \dots, a^*) + prsW_{m+1}(f(x, a, a^*, \dots, a^*)) + (1-p)rW_m(f(x, a, a^*, \dots, a^*)) \right].$$