# Coalition Formation and Price of Anarchy in Cournot Oligopolies

Nicole Immorlica[1], Evangelos Markakis[2,⋆], and Georgios Piliouras[3,⋆⋆]

[1] Northwestern University, Dept. of EECS, Evanston, IL, USA
nicimm@gmail.com
[2] Athens Univ. of Economics and Business, Dept of Informatics, Athens, Greece
markakis@gmail.com
[3] Georgia Institute of Technology, Dept. of EE, Atlanta, GA, USA
John Hopkins University, Dept. of Economics, Baltimore, MD, USA
georgios.piliouras@gmail.com

**Abstract.** Non-cooperative game theory purports that economic agents behave with little regard towards the negative externalities they impose on each other. Such behaviors generally lead to inefficient outcomes where the social welfare is bounded away from its optimal value. However, in practice, self-interested individuals explore the possibility of circumventing such negative externalities by forming coalitions. What sort of coalitions should we expect to arise? How do they affect the social welfare?

We study these questions in the setting of Cournot markets, one of the most prevalent models of firm competition. Our model of coalition formation has two dynamic aspects. First, agents choose strategically how to update the current coalition partition. Furthermore, coalitions compete repeatedly between themselves trying to minimize their long-term regret. We prove tight bounds on the social welfare, which are significantly higher than that of the Nash equilibria of the original game. Furthermore, this improvement in performance is robust across different supply-demand curves and depends only on the size of the market.

## 1 Introduction

It is a basic tenet of algorithmic game theory that agents act selfishly in the pursuit of their own interests. Borrowing from economics, the literature purports that these agents will take actions that lead to a Nash equilibrium (or a related solution concept). Hence the actions could be potentially far from the social optimum. For example, in a road network, each driver, observing traffic patterns, selects the route which minimizes his own delay. The resulting total delay can be much greater than that of the optimal flow.

⋆ Partially supported by the EU FP7 Network of Excellence Euro-NF.
⋆⋆ Supported by NSF grants CCF-0325453, AF-0910940, AFOSR grant FA9550-09-1-0420 and ONR grant N00014–09-1-0751.

A. Saberi (Ed.): WINE 2010, LNCS 6484, pp. 270–281, 2010.
© Springer-Verlag Berlin Heidelberg 2010

In a seminal paper in 1999, Koutsoupias and Papadimitriou [16] initiated the investigation of the so-called *price of anarchy* which measures the ratio of the social value in the worst-case equilibrium to the optimal social value. Recent years have seen a profusion of results exploring the price of anarchy of various non-cooperative games. The traffic example mentioned above, known as *selfish routing* in the literature, has a bounded price of anarchy of 4/3 for linear latency functions [21]. This can be viewed as a positive result. However, many settings have a drastically large price of anarchy, e.g., Cournot oligopoly games, which model competition between firms, have a linear price of anarchy for certain production functions [15].

The pursuit of self-interest however may very well encourage cooperation between agents. Such cooperation will almost certainly alter the set of stable outcomes. In an attempt to understand the implications of these issues on the price of anarchy, recent papers have studied the quality of outcomes which are stable against either all possible coalitions [2] or against arbitrary but exogenously defined coalition structures[14],[10] (e.g., the worst possible partitioning of the agents). The effect of coalition formation on social welfare has been shown to be extremely unpredictable ranging anywhere from significant improvements [2], to slight changes [10], all the way to vast degradation [14].

Tackling the issue of cooperation is pivotal in making accurate predictions about the quality of stable outcomes, especially in settings where coalitions are likely to arise. However, the theoretical models that have been introduced so far, focus mostly on the extreme cases, where either any coalition is enforceable or an arbitrary, static coalition structure is exogenously defined. Here, we introduce a model that allows for *strategically evolving* coalition structures and we examine how endogenously formed coalitions affect the quality of stable outcomes.

We focus on the setting of oligopolistic (Cournot) markets, where coalitions are known to arise in practice and we define a coalition formation game on top of the market that captures the dynamic evolution of cooperation. In our coalition formation game actions correspond to changes in the current coalition structure, hence the strategy space of the game evolves over time. Specifically, a new coalition can be created by a merger between two or more existing coalitions. An existing coalition can also be destroyed due to a deviation of a subset of its current players who decide either to form a coalition by themselves or join an existing coalition. For a new coalition to be formed, it must be the case that its creation benefits all its members.

Given a current coalition structure, we treat each coalition as a super-player who, as in [14], acts on behalf of its members and tries to maximize its aggregate utility. Any such game between the super-players (coalitions) has Nash equilibria and in the case of Cournot oligopolies we show that the utilities of the super-players at Nash equilibria are unique. This defines the value of a coalition given the current partition, which is reminiscent of the approach in [20]. Finally we divide this utility equally among the members of a coalition, since in symmetric Cournot games all players have equal production costs.

Given the rules of the game described above, we are interested in stable coalition configurations, i.e., partitions where no profitable deviating actions exist with regard to the allowed actions we have defined. We analyze the social welfare of the worst such stable partition and compare it to the cost of the optimum and refer to this ratio as the price of anarchy of our coalition formation game. We find that the price of anarchy of our coalition formation game for Cournot oligopolies is $\Theta(n^{2/5})$, where $n$ is the number of firms that participate in the market, implying a significant improvement of the actual price of anarchy of Cournot oligopolies which is $\Theta(n)$.

The value assignment to coalitions, as described in the previous paragraphs relies on the assumption that if a coalition structure is stable and hence not transient, then the super-players coalitions will reach a Nash equilibrium. We show that we can weaken this assumption considerably. Specifically, we can show that if the coalitions participate in the Cournot oligopoly repeatedly in a fashion that minimizes their long term regret then the average utility of the super-players (coalitions) will converge to their levels at Nash equilibria. Regret compares the average utility of a player to that of the best fixed constant action with hindsight. Having no-regret means that no deviating action would significantly improve the firm's utility. Several learning algorithms are known to provide such guarantees ([3,23] and references therein). More importantly, the assumption is not tied to any specific algorithmic procedure, but instead captures successful long-term behavior. Finally, since the setting of oligopolies markets is in its nature repeated, this observation significantly strengthens the justification of our model.

**Paper Structure.** Section 2 offers the definition of Cournot oligopolies and a detailed exposition of our coalition formation model. In Section 3 we prove tight bounds for the price of anarchy of the Cournot coalition formation game. Finally, Section 4 extends our analysis to the case of no-regret behavior.

### 1.1 Related Work

Quantifying the inefficiency of outcomes when coalitions are allowed to form has been the subject of much recent work. In [14], the authors initiate the study of the *price of collusion*, which is a measure of the inefficiency of the worst possible partition of the set of players. In [10], both the quality and tractability of stable outcomes is examined in atomic congestion games with coalitions. The models above do not raise any strategic issues in the formation of the coalitions, which are essentially exogenously enforced upon the game. In contrast, we focus on a strategic setting, where we study only stable partitions, i.e., partitions where agents have no incentive to deviate, as we define in Section 2.

Other notions of inefficiency have also been analyzed. In [2], the authors analyze the *price of strong anarchy*, i.e., the inefficiency of Nash equilibria which are resilient to deviations by coalitions[1]. In [5], a different measure is introduced, namely the *price of democracy*. This notion captures the inefficiency of a given

[1] Unfortunately, for several classes of games including Cournot markets, strong Nash equilibria do not exist.

coalition formation process (e.g. a bargaining process) with respect to a cooperative game. The authors study this notion in the context of weighted voting games for certain intuitive bargaining processes. Hence the inefficiency is measured with regard to the arising partitions in the subgame perfect equilibria of the corresponding bargaining game.

Regarding Cournot games, it has been long known that the loss of efficiency at Nash equilibria can be quite high. Earlier studies focused on empirical analysis [12] whereas more recently, price of anarchy bounds have been obtained in [11,15]. Collusion and cartel enforcement in Cournot games have been studied experimentally, see e.g., [22]. Mechanism design aspects of collusion have also been explored, see [7]. For more on Cournot games and their variants, we refer the reader to [17].

Dynamic coalition formation has been studied extensively both in the economics as well as in the computer science literature. We refer the reader to [8,4,13,19] and [1][Section 5.1] as well as the numerous references therein. The main goals of these works have been to provide appropriate game theoretic solution concepts (both from a cooperative and noncooperative point of view) and to design intuitive procedures that converge experimentally or theoretically to such solution concepts.

Conceptually, the closest example to our approach that we know of, is the work of Ray and Vohra in [20]. The authors propose a solution concept ("binding agreement") that allows for the formation of coalition structures and examine the inefficiency of stable partitions. Unlike in our work, their deviations can only make the existing coalition structure finer- never coarser. In the case of symmetric Cournot games, it is shown that there always exists a stable partition with social welfare $O(\sqrt{n})$ worse than the optimal. However, the social welfare of the worst stable partition is always at least as bad as that of the worst Nash. In follow-up work [19], Ray analyzes a class of bargaining processes which assumes players with infinite foresight and shows that in symmetric Cournot games the only coalition structure that is stable for all of them, has social welfare $\Theta(\sqrt{n})$ worse than the optimal.

Finally, there has been some recent work on the behavior of no-regret algorithms in Cournot oligopolies. In [9], [18] several convergence results are shown for different classes of Cournot oligopolies. To our knowledge our paper is the first to consider the behavior of *coalitions* which are behaving in a no-regret fashion in any kind of setting.

## 2 The Model

We will demonstrate the main point of our work in the context of Cournot games. The definitions presented in this Section can be easily generalized and applied to other contexts but we postpone a more general treatment for an extended version.

Cournot games describe a fundamental model of competition between firms. They were introduced by Cournot in his much celebrated work [6]. In Cournot games, firms control their production levels and by doing so influence the market prices. In the simplest Cournot model all the firms produce the same good; the demand for this product is linear in the total production (i.e. the price decreases linearly with total production); the unit cost of production is fixed and equal across all firms. The revenue of a firm is the product of the firm's part of the market production times the price. Finally, the utility of a firm is equal to its revenue minus its total production cost. Overproducing leads to low prices, while at the same time an overly cautious production rate leads to a small market share and reduced revenue. The balancing act between these two competing tendencies is known to give rise to a unique Nash equilibrium. More formally:

**Definition 1.** *A linear and symmetric Cournot oligopoly is a noncooperative game between a set $N = \{1, 2, ..., n\}$ of players (firms), all capable of producing the same product. The strategy space of each firm is $\mathbb{R}_+$, corresponding to the quantity of the product that the firm decides to produce. Given a profile of strategies, $q = (q_1, ..., q_n)$, the utility of firm $i$ is $u_i(q) = q_i p(q) - cq_i$, where $p(q)$ is the price of the product, determined by $p(q) = \max\{0, a - b \sum_i q_i\}$, for some parameters $a, b$, and $c$ is a production cost, with $a > c$.*

**Proposition 1 ([6]).** *In the unique Nash equilibrium of a Cournot oligopoly with $n$ players, the production level is the same for all players and equal to $q_i = q^* = \frac{(a-c)}{b(n+1)}$. The utility of each player is equal to $u_i = \frac{(a-c)^2}{b(n+1)^2}$ and the social welfare is equal to $\frac{(a-c)^2 n}{b(n+1)^2}$.*

### 2.1 Cournot Games with a Fixed Partitioning of the Players

Suppose now that the players are given the opportunity to form coalitions and sign agreements with other firms, as a means of reducing competition and improving on their welfare. Given a partition of the players into coalitions, we can think of the new situation as a super-game whose super-players are the coalitions themselves. The strategy for a coalition, or super-player, is now a vector assigning a strategy to each of its members. The payoff to the super-player is the aggregate payoff its members would achieve with their assigned strategies in the original game. This definition can be used to model coalitions in general games as in [10,14].

**Definition 2.** *Let $\mathcal{G}$ be a game of $n$ players, with $A_j$ being the set of available actions and $u_j^{\mathcal{G}}(a_1, \ldots, a_n)$ the utility function for each player $j$. Given a partitioning $\Pi = (S_1, ..., S_k)$ of the players, then the corresponding super-game consists of the following:*

- *$k$ super-players*
- *The strategy set for super-player $S_i$ is the set of vectors $\overrightarrow{a}_{S_i} \in \prod_{j \in S_i} A_j$.*
- *The utility of super-player $S_i$ is $u_{S_i}(\overrightarrow{a}_{S_1}, \ldots, \overrightarrow{a}_{S_k}) = \sum_{j \in S_i} u_j^{\mathcal{G}}(a_1, \ldots, a_n)$ where $a_j$ is the strategy assigned to player $j$ by his coalition $S_i$ in the coalition's strategy $\overrightarrow{a}_{S_i}$.*

It is straightforward to check that for Cournot games, the super-game with $k$ super-players is essentially equivalent to a Cournot game with $k$ players[2].

**Lemma 1.** *Consider a Cournot oligopoly super-game for a fixed partitioning $\Pi = (S_1, ..., S_k)$ of players. The players' utilities and the social welfare in this game under any strategy profile $\overrightarrow{q}_{S_1}, \ldots, \overrightarrow{q}_{S_k}$ (where $\overrightarrow{q}_{S_i} \in \mathcal{R}_+^{|S_i|}$) are equal to the corresponding utilities and social welfare of a linear and symmetric Cournot game with k players where each player i produces the aggregate production $\sum_{j \in S_i} (\overrightarrow{q}_{S_i})_j$ of the corresponding coalition $S_i$. Furthermore, a strategy profile for the super-game with the fixed partitioning is a Nash equilibrium if and only if the k-tuple of the aggregate levels of productions for each coalition is the unique Nash equilibrium for the Cournot game on k players (without coalitions).*

Lemma 1 allows us to use theorems regarding Cournot games to study the Nash equilibria and welfare of Cournot games with coalitions. Specifically, it implies that the social welfare is the same in all Nash equilibria of the Cournot game with a fixed partitioning. Hence we can define the price of anarchy as the ratio of this social welfare over the optimal social welfare, which is realized when all agents unite into a single coalition. By combining proposition 1 with lemma 1 we derive:

**Lemma 2.** *The price of anarchy of a Cournot oligopoly with a fixed partition $\Pi = (S_1, ..., S_k)$ is $\frac{(k+1)^2}{4k}$.*

As a consequence, the price of anarchy in the original noncooperative Cournot oligopoly with $n$ players is very high, namely linear in the number of players, as has been observed previously [15].

**Corollary 1.** *The price of anarchy in the original Cournot game with n players, where no coalitions are allowed to form is $\Theta(n)$.*

### 2.2 Cournot Coalition Formation Games

Next, we move away from the fixed coalition structure assumption and instead we will allow the players to dynamically form coalitions. We will call this game the *Cournot coalition formation game.* Given some initial partition, players or sets of players can consider deviations according to the rules that we define below. As we have seem by Lemma 1, for any resulting partition, say with $k$ coalitions, the utility of each coalition is unique in all Nash equilibria of the Cournot game with fixed coalitions, and equal to the utility of a player in the unique Nash equilibrium of a symmetric Cournot game with $k$ players. In the coalition formation game, each of the $n$ players, when evaluating a possible action of hers, estimates her resulting utility to be equal to her equiproportional share of the Nash equilibrium utility of the coalition to which she belongs, given the resulting coalition structure. More formally:

[2] Henceforth, when it is clear from the context, we will use game instead of super-game and player instead of super-player.

**Definition 3.** *We define a coalition formation game on top of a symmetric Cournot game to consist of the following:*

- *$n$ players and a current partitioning of them into $k$ coalitions $\Pi = (S_1, ..., S_k)$*
- *Given the current partition $\Pi$, the allowed moves (deviations) that players can use along with the consequences for the coalition left behind (i.e., the non-deviators) are as follows:*

  Type 1: *A subset $S'_i$ of a current coalition $S_i$ decides to deviate and form a new coalition. The rest of the members, if any, of the original coalition (i.e. $S_i/S'_i$) dissolve into singletons.*[3]

  Type 2: *A strict subset $S'_i$ of a current coalition $S_i$ decides to leave its current coalition $S_i$ and join another coalition of $\Pi$, say $S_j$. The rest of the members of the original coalition (i.e. $S_i/S'_i$) dissolve into singletons.*

  Type 3: *A set of coalitions of $\Pi$ decide to unite and form a coalition. The rest of the coalitions remain as they were.*
- *Given a partition $\Pi$ and a player $i$ in coalition $S_j$ of $\Pi$, denote by $u_{S_j}(\Pi)$, the uniquely defined Nash utility of coalition $S_j$ in the symmetric Cournot game with fixed coalition structure $\Pi$. The utility of player $i$ in this case, is defined to be equal to $u_{S_j}(\Pi)/|S_j|$.*

In terms of our assumptions about allowable actions, unlike the work of [20], we allow both the creation as well as the destruction of coalitions. Furthermore, we assume that in some of the deviating actions (Type 1 and 2), the leftover coalition from where the deviation emerged, dissolves into singletons. This is reminiscent of past approaches[7,13]. Essentially, our assumption encodes that non-deviators will react cautiously.

We will be interested in analyzing the price of anarchy for partitions in which no player or set of players has an incentive to change the current coalition structure. In order to characterize stable coalitions, we need to define when a deviation is successful. A deviation is successful if and only if the utility of all the players that induce this deviation strictly increases as a result. More formally:

**Definition 4.** *A deviation is successful iff all the players that facilitate the deviation strictly increase their payoff by doing so. Specifically, a deviation of*

- Type 1 *is successful iff all the players in $S'_i$ increase their payoffs.*
- Type 2 *is successful iff all the deviating players in $S'_i$ as well as all the members of the coalition $S_j$ who accept them increase their payoffs.*
- Type 3 *is successful iff all the members of all the merging coalitions increase their payoffs.*

**Definition 5.** *A partition $\Pi$ is stable if there exists no successful deviation of any type.*

[3] This type of actions also includes the non-action option (i.e. the coalition structure remains unaltered), when $S'_i = S_i$.

In the usual manner of the "price of anarchy" literature, we are interested in bounding the ratio of the social welfare of the worst stable outcome (i.e. coalition partition) divided by the optimal social welfare. In our setting, the stable outcomes do not correspond exactly to Nash, since we allow bilateral moves (e.g. type 3). Nevertheless, we will still use the term price of anarchy to refer to this ratio, since it characterizes the loss in performance due to the lack of a centralized authority that could enforce the optimal (grand) coalition.

**Definition 6.** *Given a Cournot coalition formation game, we define the price of anarchy as the ratio of the social welfare that is achieved at the worst stable partition divided by the optimal social welfare.*

# 3 The Main Result

The starting point of our work is the observation of Corollary 1 that without coalition formation the price of anarchy is $\Theta(n)$. Hence our goal is to understand the quality of the worst stable partition structure and compare it to the optimal. The optimal partition structure is trivially the one where all players have united in a single coalition, as there is no competition in such a setting. Our main result is that the price of anarchy is significantly reduced when coalition formation is allowed. Formally:

**Theorem 1.** *The price of anarchy of the coalition formation game is $\Theta(n^{2/5})$.*

## 3.1 The Proof of the Upper Bound

We begin by proving that the price of anarchy is $O(n^{2/5})$. We will first establish this upper bound on a restricted version of our model. In particular, we restrict each type of the allowed deviations of Definition 3 as follows:
*Type 1:* A member of a coalition of $\Pi$, decides to form a singleton coalition on his own. The coalition from which the player left dissolves into singleton players.
*Type 2:* A member of a coalition of $\Pi$ decides to leave its current coalition $S_i$ (where $|S_i| \geq 2$), and join another coalition of $\Pi$, say $S_j$. The rest of coalition $S_i$ dissolves into singleton players.
*Type 3:* A set of singleton players of $\Pi$ decide to unite and form a coalition.

We will refer to this game as the *restricted coalition formation game*. Once we establish the upper bound in the restricted model, it is trivial to extend it to the general model since the set of stable partitions only gets smaller in the general model. To analyze the price of anarchy, we will derive a characterization of the stable partitions. Throughout the analysis we will normally denote the cardinality of a coalition $S_i$ by $s_i = |S_i|$. The first Lemma below says that for coalitions of size at least 2, its members need only consider Type 1 deviations.

**Lemma 3.** *Consider a partition $\Pi = (S_1, ..., S_k)$, with $k \geq 2$. For a player that belongs to a coalition of $\Pi$ of size at least 2, the most profitable deviation (though not necessarily a successful one) is the deviation where the player forms a singleton coalition on his own.*

*Proof.* Consider a coalition $S_i$ of $\Pi$ of size $s_i$. Suppose $s_i \geq 2$ and consider a player $j \in S_i$. The available deviations for $j$ are either to form a coalition on his own or to join an existing coalition. In the former case, the coalition $S_i$ will dissolve and the total number of coalitions in the new game will be $k + s_i - 1$. Hence the payoff of $j$ will be $u = \frac{(a-c)^2}{b(k+s_i)^2}$. On the other hand, if $j$ goes to an existing coalition, then $S_i$ again dissolves but the total number of coalitions is now $k + s_i - 2$. Since $j$ will be in a coalition with at least 2 members, the payoff to $j$ will be at most: $u' \leq \frac{(a-c)^2}{2b(k+s_i-1)^2}$.

We wish to have $u \geq u'$. It suffices to show that $(k + s_i)^2 \leq 2(k + s_i - 1)^2$, which is equivalent to $(k + s_i)^2 - 4(k + s_i) + 2 \geq 0$. For this it suffices to show that $(k + s_i) \geq 2 + \sqrt{2}$. But we have assumed that $k \geq 2$ and that $s_i \geq 2$, hence the proof is complete. □

The next lemma is based on Lemma 3 and characterizes coalitions of size at least 2, for which there are no successful deviations for its members.

**Lemma 4.** *Consider a partition $\Pi = (S_1, ..., S_k)$, with $k \geq 2$. For a coalition $S_i$ with $s_i \geq 2$, there is no successful deviation for its members iff $s_i \geq k^2$.*

*Proof.* Consider a coalition of partition $\Pi$, say $S_i$ with $s_i \geq 2$. The payoff that a player in $S_i$ now receives is $u = \frac{(a-c)^2}{s_i b(k+1)^2}$. By Lemma 3 the most profitable deviation for any player of $S_i$ is to form a singleton coalition, in which case he would receive a payoff of $u = \frac{(a-c)^2}{b(k+s_i)^2}$. In order that no player has an incentive to deviate, we need that $(k + s_i)^2 \geq s_i(k + 1)^2$, which is equivalent to $s_i \geq k^2$. □

We now deal with deviations of players that form singleton coalitions in a partition $\Pi$. By definition, we only need to consider Type 3 deviations for singleton players.

**Lemma 5.** *Consider a partition $\Pi = (S_1, ..., S_k)$, with $k \geq 2$. Suppose that $\Pi$ contains $k_1$ singleton coalitions with $k_1 \geq 2$, and $k_2$ non-singleton ones ($k_1+k_2 = k$). The merge of the $k_1$ singletons is not a successful deviation iff $k_1 \leq (k_2+1)^2$.*

*Proof.* The $k_1$ singletons receive in $\Pi$ a payoff of $(a-c)^2/(b(k_1+k_2+1)^2)$. After the merge, their payoff will be $(a - c)^2/(k_1 b(k_2 + 2)^2)$. Hence, the merge will not be successful, iff $(a - c)^2/(b(k_1 + k_2 + 1)^2) \geq (a - c)^2/(k_1 b(k_2 + 2)^2)$. By manipulation of terms this is shown to equivalent to $(k_2 + 1)^2 \geq k_1$. □

Finally we show that for ensuring stability there is no need to consider any other Type 3 deviation of smaller coalitions.

**Lemma 6.** *Consider a partition $\Pi = (S_1, ..., S_k)$, with $k \geq 2$ and suppose that it contains $k_1$ singleton coalitions with $k_1 \geq 2$, and $k_2$ non-singleton ones. There is a successful Type 3 deviation iff the merge of all $k_1$ singletons is a successful deviation.*

*Proof.* One direction is trivial, namely if the merge of all $k_1$ singletons is a successful deviation. For the reverse direction, suppose there is a successful type

3 deviation which is not the merge of all the $k_1$ singletons. Let $m$ be the number of players who merge and suppose $2 \leq m < k_1$. By arguing as in Lemma 5, we get that in order for the deviation to be successful, it should hold that $(k_1+k_2+1)^2 > m(k_1+k_2-m+2)^2$. Let $\lambda = k_2+1$ and $\theta = \lambda + k_1 - m$. Restating the condition in terms of $\lambda$ and $\theta$ we get $(k_1+\lambda)^2 = (\theta+m)^2 > m(\theta+1)^2$, which via a rearranging of terms can be shown to be equivalent to $m > \theta^2$.

However, we have that $k_1 > m > (\lambda+k_1-m)^2 > \lambda^2 = (k_2+1)^2$. By Lemma 5, this means that the merge of all $k_1$ singletons is also a successful deviation. □

All the above can be summarized as follows:

**Corollary 2.** *Consider a partition $\Pi$. For $n \leq 2$, $\Pi$ is stable iff it is the grand coalition. For $n \geq 3$, suppose $\Pi = (S_1, ..., S_k)$ with $k_1$ singleton coalitions and $k_2$ non-singleton ones. Then $\Pi$ is stable iff it is either the grand coalition or the following hold:*

- $k_1 \leq (k_2+1)^2$.
- *For every non-singleton coalition $S_i$, $s_i \geq k^2$.*

Having acquired a characterization of the stable partitions, we can now analyze the (pure) price of anarchy of the restricted coalition formation game on top of a symmetric Cournot oligopoly. We omit the proof due to lack of space.

**Theorem 2.** *The price of anarchy of the restricted coalition formation game under symmetric Cournot oligopoly is $O(n^{2/5})$, where $n$ is the total number of players.*

Finally, we come back to the original coalition formation game of Definition 3. Since in that game we have only enlarged the set of possible deviations with regard to the restricted coalition formation game, the set of stable partitions can only decrease. As a result, the price of anarchy for the original game is also $O(n^{2/5})$. This completes the proof for the upper bound of Theorem 1.

### 3.2 The Construction of the Lower Bound

The lower bound is obtained by the construction in the following lemma whose proof appears in the extended version of our paper:

**Lemma 7.** *For any $N$, let $n$ be the number: $n = \lceil 4N^{4/5} \rceil \lfloor N^{1/5} \rfloor + \lfloor N^{2/5} \rfloor$. Consider a game on $n$ players and a partition of the $n$ players consisting of $k_1 = \lfloor N^{2/5} \rfloor$ singletons and $k_2 = \lfloor N^{1/5} \rfloor$ coalitions of size $s = \lceil 4N^{4/5} \rceil$ each. This coalition structure is stable for the Cournot coalition formation game.*

Since the total number of coalitions in the construction is $k = k_1 + k_2 \geq N^{2/5} = \Omega(n^{2/5})$, by Lemma 2, we obtain the desired lower bound.

**Theorem 3.** *For any number of players $n$, there exist stable partitions with cost $\Omega(n^{\frac{2}{5}})$ the cost of the optimal partition.*

## 4 Coalition Formation under No-Regret

So far, given a partition $\Pi$, we assign to each coalition $S_i$ value equal to its uniquely defined utility at the Nash equilibria of the Cournot game with a fixed coalition partition $\Pi$. The reasoning behind this is that if the coalition partition $\Pi$ is not transient, then the players/coalitions will hopefully reach a Nash equilibrium and hence their uniquely defined utility at it, is a good estimator of how much they value their current coalition partition.

Here, we argue that we can significantly weaken the assumption that the players/coalitions will reach an equilibrium. In fact, we will establish that if the coalitions participate in the Cournot oligopoly repeatedly in a fashion that minimizes their long term regret then their average utility will converge to their levels at Nash equilibria. The regret of an online learning algorithm[4] is defined as the maximum over all input instances of the expected difference in payoff between the algorithms actions and the best action. If this difference is guaranteed to grow sublinearly with time, we say it is a no-regret learning algorithm [3,23].

This notion captures successful long-term behavior and can be achieved in practice by several natural learning algorithms [3,23]. Putting all these together, we have that the values assigned to coalitions by our model, are in excellent agreement with the average utilities they would actually receive by participating repeatedly (and successfully) in the market.

**Theorem 4.** *Consider a Cournot oligopoly game with a fixed partitioning of the $n$ players in $k$ coalitions $\Pi = (S_1, ..., S_k)$. If all $k$ (super-)players employ no-regret strategies, then their average utilities converge to their Nash levels.*

## 5 Discussion and Future Work

We have introduced a model of coalition formation and have identified tight bounds on the inefficiency of stable coalitions in oligopolistic markets. Our approach combines elements of cooperative game theory (e.g. payoff distribution amongst the members of the coalition) and noncooperative game theory (e.g. coalitions compete against each other). Such balancing acts of cooperation-competition are common in real-life economic settings and we believe that this work opens up a promising avenue for future research. The natural next step would be to examine the sensitivity of our results to changes in the underlying coalition formation process, as well as extensions to different classes of games.

## References

1. Airiau, S.: Lecture Notes in Coalitional Games. In: 11th European Agent Systems Summer School, EASSS 2009 (2009), `http://staff.science.uva.nl/~stephane/Teaching/EASSS09/notes-coopGT-easss09.pdf`

[4] An online learning algorithm is an algorithm for choosing a sequence of elements of some fixed set of actions, in response to an observed sequence of cost functions mapping actions to real numbers. The $t$-th action chosen by the algorithm may depend on the first $t-1$ observations but not on any later observations.

2. Andelman, N., Feldman, M., Mansour, Y.: Strong price of anarchy. In: ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 189–198 (2007)
3. Cesa-Bianchi, N., Lugosi, G.: Prediction, Learning and Games. Cambridge University Press, Cambridge (2006)
4. Chalkiadakis, G.: A Bayesian Approach to Multiagent Reinforcement Learning and Coalition Formation under Uncertainty. PhD thesis, Department of Computer Science, University of Toronto, Canada (2007)
5. Chalkiadakis, G., Elkind, E., Polukarov, M., Jennings, N.R.: The price of democracy in coalition formation. In: International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS), pp. 401–408 (2009)
6. Cournot, A.A.: Recherches sur les Principes Mathématiques de la Théorie des Richesses. Hatchette, Paris (1838); English translation: Researches into the Mathematical Principles of the Theory of Wealth. Macmillan, New York (1897)
7. Cramton, P., Palfrey, T.: Cartel enforcement with uncertainty about costs. International Economic Review 31(1), 17–47 (1990)
8. Dieckmann, T., Schwalbe, U.: Dynamic Coalition Formation and the Core (1998); Economics Department Working Paper Series, Department of Economics, National University of Ireland - Maynooth
9. Even-Dar, E., Mansour, Y., Nadav, U.: On the convergence of regret minimization dynamics in concave games. In: ACM Symposium on Theory of Computing, STOC (2009)
10. Fotakis, D., Kontogiannis, S., Spirakis, P.: Atomic congestion games among coalitions. ACM Trans. Algorithms 4(4), 1–27 (2008)
11. Guo, X., Yang, H.: The price of anarchy of cournot oligopoly. In: Deng, X., Ye, Y. (eds.) WINE 2005. LNCS, vol. 3828, pp. 246–257. Springer, Heidelberg (2005)
12. Harberger, A.: Monopoly and resource allocation. Amer. Econ. Review 44, 77–87 (1954)
13. Hart, S., Kurz, M.: Endogenous formation of coalitions. Econometrica 51, 1047–1064 (1983)
14. Hayrapetyan, A., Tardos, É., Wexler, T.: The effect of collusion in congestion games. In: ACM Symposium on Theory of Computing (STOC), pp. 89–98 (2006)
15. Kluberg, J., Perakis, G.: Generalized quantity competition for multiple products and loss of efficiency. In: 46th Annual Allerton Conference on Communication, Control, and Computing, pp. 930–936 (2008)
16. Koutsoupias, E., Papadimitriou, C.: Worst-case equilibria. Computer Science Review 3(2), 65–69 (2009)
17. Mas-Colell, A., Green, J., Whinston, M.D.: Microeconomic Theory. Oxford University Press, Oxford (1995)
18. Nadav, U., Piliouras, G.: No regret learning in oligopolies: Cournot vs bertrand. In: 3rd International Symposium on Algorithmic Game Theory, SAGT (2010)
19. Ray, D.: A Game-Theoretic Perspective on Coalition Formation. Oxford Univ. Press, Oxford (2007)
20. Ray, D., Vohra, R.: Equilibrium binding agreements. Journal of Economic Theory 73, 30–78 (1997)
21. Roughgarden, T., Tardos, E.: How bad is selfish routing? J. ACM 49(2), 236–259 (2002)
22. Suetens, S., Potters, J.: Bertrand colludes more than Cournot. Experimental Economics 10, 71–77 (2007)
23. Young, H.: Strategic Learning and Its Limits. Oxford University Press, Oxford (2004)