

Coalition Formation and Price of Anarchy in Cournot Oligopolies

Nicole Immorlica* Evangelos Markakis† Georgios Piliouras‡

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.

*Northwestern University, Dept. of EECS, Evanston, IL, USA. Email: nicimm@gmail.com.

†Athens Univ. of Economics and Business, Dept of Informatics, Athens, Greece.

Email: markakis@gmail.com. Partially supported by the Euro-NF Network of Excellence.

‡Georgia Institute of Technology, Dept. of EE, Atlanta, GA, USA. Johns Hopkins University, Dept. of Economics, Baltimore, MD, USA. Email:georgios.piliouras@gmail.com.

Supported by NSF grants CCF-0325453, AF-0910940, AFOSR grant FA9550-09-1-0420 and ONR grant N00014-09-1-0751.

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 optimal. 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.

In a seminal paper in 1999, Koutsoupias and Papadimitriou [18] 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 [23]. 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 [16].

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 [15],[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 [15].

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 [15], 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 [22]. 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 of the underlying game. 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, 25] 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 as well as 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 [15], 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 coalitions¹. In [5], a different measure is introduced, namely the *price of democracy*. This notion captures the inefficiency of a given 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

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

can be quite high. Earlier studies focused on empirical analysis [12] whereas more recently, price of anarchy bounds have been obtained in [11, 16]. Collusion and cartel enforcement in Cournot games have been studied experimentally, see e.g., [24]. Mechanism design aspects of collusion have also been explored, see [7].

Dynamic coalition formation has been studied extensively both in the economics as well as in the computer science literature. We refer the reader to [4, 13, 14, 21] 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 [22]. 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 [21], 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], [20] 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 the 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 $N = \{1, 2, \dots, n\}$ 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, $q = (q_1, \dots, q_n)$, 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 $u_i = \frac{(a-c)^2}{b(n+1)^2}$.*

For more on Cournot games and their variants, we refer the reader to [19].

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, 15].

Definition 2 *Let \mathcal{G} be a game of n players, with A_j being the set of available actions and $u_j^{\mathcal{G}}(a_1, \dots, a_n)$ the utility function for each player j . Given a partitioning $\Pi = (S_1, \dots, S_k)$ 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 $\vec{a}_{S_i} \in \prod_{j \in S_i} A_j$.
- The utility of super-player S_i is given by $u_{S_i}(\vec{a}_{S_1}, \dots, \vec{a}_{S_k}) = \sum_{j \in S_i} u_j^{\mathcal{G}}(a_1, \dots, a_n)$ where a_j is the strategy assigned to player j by his coalition S_i in the coalition's strategy \vec{a}_{S_i} .

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

The following Lemma demonstrates that for Cournot games, the super-game with k super-players is essentially equivalent to a Cournot game with k players. This allows us to use theorems regarding Cournot games to study the Nash equilibria and welfare of Cournot games with coalitions. The proof appears in Appendix A.

Lemma 2 *Consider a Cournot oligopoly super-game for a fixed partitioning $\Pi = (S_1, \dots, S_k)$ of players. The social welfare in this game under any strategy profile $\vec{q}_{S_1}, \dots, \vec{q}_{S_k}$ (where $\vec{q}_{S_i} \in \mathcal{R}_+^{|S_i|}$) is equivalent to the social welfare of a linear and symmetric Cournot game with k players where each player i produces the aggregate production $\sum_{j \in S_i} (\vec{q}_{S_i})_j$ 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 2 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. The following Lemma shows that the price of anarchy depends only on the number of coalitions in the partition. The proof appears in Appendix B.

Lemma 3 *The price of anarchy of a Cournot oligopoly with a fixed partition $\Pi = (S_1, \dots, S_k)$ 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 [16].

Corollary 4 *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 seen by Lemma 2, 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:

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 $\Pi = (S_1, \dots, S_k)$,
- Given the current partition Π , 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. ²
 - 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 Π , 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 Π decide to unite and form a coalition. The rest of the coalitions remain as they were.

²This type of actions also includes the non-action option (i.e. the coalition structure remains unaltered), when $S'_i = S_i$.

- Given a partition Π and a player i in coalition S_j of Π , denote by $u_{S_j}(\Pi)$, the uniquely defined utility of coalition S_j in the symmetric Cournot game with fixed coalition structure Π . The utility of player i in this case, is defined to be equal to $u_{S_j}(\Pi)/|S_j|$.

Here, we will discuss the assumptions behind our model. In terms of allowable actions, unlike the work of [22], we allow both the creation as well as the destruction of coalitions. Furthermore, we not only allow for the creation of new coalitions but we put no restrictions on the number of merging partners, resulting in a more expressive model. Furthermore, we assume that in some of the deviating actions (Type 1 and 2), the leftover coalition from where the deviation emerged, dissolves into singleton players. This assumption has been considered before in other works on Cournot games, such as [7], (where cartel enforcement is studied in Cournot games from a mechanism design point of view) but it has also been considered more generally in the context of defining models of coalition formation for arbitrary games. As an example, Hart and Kurz [14] define and study stability notions for cooperative games with coalition structures. One of the main concept that they consider is what they call γ -stability which essentially corresponds to a strong equilibrium under the assumption that coalitions which the deviators abandoned are breaking up into singletons. A related work by the same authors with more examples appears in [13]. Other behavioral models for the non-deviators have also been proposed in the context of partition function games, where the value of a coalition depends on how the remaining people react to a deviation, see e.g. [17][Section 3.1] for a summary on this. But clearly, as pointed out in [14], there is no universally correct answer on modeling how the non-deviators react. We think however that for Cournot games that capture competition among firms, our model is a reasonable approach since, once there is a deviation, the left over firms may not be able to sustain a commercial agreement and may have to dissolve before reconsidering the terms of continuing their cooperation. Essentially we assume that the players resort to an individual best reply action, since their new payoff in the new partition is the payoff at a Nash equilibrium with these players as singletons.

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, we have that:

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 Π is stable if there exists no successful deviation of any type.*

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 4 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 setting. Our main result is that the price of anarchy is significantly reduced when coalition formation is allowed. Formally:

Theorem 5 *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 Π , 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 Π decides to leave its current coalition S_i (where $|S_i| \geq 2$), and join another coalition of Π , say S_j . The rest of coalition S_i dissolves into singleton players.

Type 3: A set of singleton players of Π 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. The reason for choosing this restricted version is that stability against this restricted set of deviations is sufficient to impose the desired upper bound as we exhibit below.

To analyze the price of anarchy, we first derive a characterization of stable partitions in the restricted game. This will depend on a number of lemmas. The bound of $n^{2/5}$ on the price of anarchy is finally derived as a solution to a certain mathematical program, which is determined by the constraints that are imposed by stability against deviations (see Theorem 11).

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 6 Consider a partition $\Pi = (S_1, \dots, S_k)$, with $k \geq 2$. For a player that belongs to a coalition of Π 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 Π 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 want to prove $u \geq u'$. 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. \square

The next lemma is based on Lemma 6 and characterizes coalitions of size at least 2, for which there are no successful deviations for its members.

Lemma 7 Consider a partition $\Pi = (S_1, \dots, S_k)$, 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 Π , 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 6 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 $u = \frac{(a-c)^2}{b(k+s_i)^2}$. Hence 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 $s_i \geq k^2$. \square

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

Lemma 8 Consider a partition $\Pi = (S_1, \dots, S_k)$, with $k \geq 2$. Suppose that Π 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 Π 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 $(a-c)^2/(b(k_1+k_2+1)^2) \geq (a-c)^2/(k_1 b(k_2+2)^2)$. As a result, we have that:

$$\begin{aligned}
(a-c)^2/(b(k_1+k_2+1)^2) &\geq (a-c)^2/(k_1b(k_2+2)^2) \Leftrightarrow \\
k_1(k_2+2)^2 &\geq (k_1+k_2+1)^2 \Leftrightarrow \\
k_1(k_2+1)^2 + 2k_1(k_2+1) + k_1 &\geq k_1^2 + (k_2+1)^2 + 2k_1(k_2+1) \Leftrightarrow \\
(k_1-1)(k_2+1)^2 &\geq k_1^2 - k_1 \Leftrightarrow \\
(k_2+1)^2 &\geq k_1
\end{aligned}$$

□

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

Lemma 9 *Consider a partition $\Pi = (S_1, \dots, S_k)$, 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 8, 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 λ and θ we get:

$$\begin{aligned}
(k_1 + \lambda)^2 = (\theta + m)^2 &> m(\theta + 1)^2 \Leftrightarrow \\
\theta^2 + m^2 + 2\theta m &> m\theta^2 + 2\theta m + m \Leftrightarrow \\
m &> \theta^2
\end{aligned}$$

However, $k_1 > m > (\lambda + k_1 - m)^2 > \lambda^2 = (k_2 + 1)^2$. By Lemma 8, this means that the merge of all k_1 singletons is also a successful deviation and the proof is complete. □

All the above can be summarized as follows:

Corollary 10 *Consider a partition Π . For $n \leq 2$, Π is stable iff it is the grand coalition. For $n \geq 3$, suppose $\Pi = (S_1, \dots, S_k)$ with k_1 singleton coalitions and k_2 non-singleton ones. Then Π 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 are ready to analyze the (pure) price of anarchy of the restricted coalition formation game on top of a symmetric Cournot oligopoly. The proof is given in appendix C.

Theorem 11 *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 the game we are interested in, 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 5.

3.2 The Proof of the Lower Bound

The lower bound is obtained by the construction in the following Lemma:

Lemma 12 *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.*

The proof of this lemma is given in Appendix D, where we show that none of the three types of deviating moves is successful for this coalition partition. 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 3, we obtain the desired lower bound for any number of players n .

Corollary 13 *For any number of players n , there exist stable partitions with cost $\Omega(n^{2/5})$ the cost of the optimal partition.*

4 Coalition Formation under No-regret

So far, given a partition Π , 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 Π . The reasoning behind this is that if the coalition partition Π 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.

In this paragraph, we show that we can significantly weaken the assumption that the players/coalitions will reach an equilibrium. In fact, we will show 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. Intuitively, having no-regret³ means that no deviating action would significantly improve the utility of the player. This notion captures successful long-term behavior and can be achieved in practice by several natural learning algorithms ([3, 25] and references therein). 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. Since the setting of oligopolies markets is in its nature repeated, this observation significantly strengthens the justification of our model.

Theorem 14 *Consider a Cournot oligopoly game with a fixed coalition partitioning of the n players in k coalitions $\Pi = (S_1, \dots, S_k)$. If all k (super-)players employ no-regret strategies, then their average utilities converge to their Nash levels.*

³Formal definitions are in appendix E along with proof of theorem 14

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] S. Airiau. Lecture Notes in Coalitional Games. Tutorial notes of the 11th European Agent Systems Summer School (EASSS-2009). Available at <http://staff.science.uva.nl/~stephane/Teaching/EASSS09/notes-coopGT-easss09.pdf>, 2009.
- [2] N. Andelman, M. Feldman, and Y. Mansour. Strong price of anarchy. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 189–198, 2007.
- [3] N. Cesa-Bianchi and G. Lugosi. *Prediction, Learning and Games*. Cambridge University Press, 2006.
- [4] G. Chalkiadakis. *A Bayesian Approach to Multiagent Reinforcement Learning and Coalition Formation under Uncertainty*. PhD thesis, Department of Computer Science, University of Toronto, Canada, 2007.
- [5] G. Chalkiadakis, E. Elkind, M. Polukarov, and N. R. Jennings. The price of democracy in coalition formation. In *International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 401–408, 2009.
- [6] A. A. Cournot. *Recherches sur les Principes Mathématiques de la Théorie des Richesses*. Paris:Hachette, 1838. English translation: *Researches into the Mathematical Principles of the Theory of Wealth*, New York: Macmillan, 1897.
- [7] P. Cramton and T. Palfrey. Cartel enforcement with uncertainty about costs. *International Economic Review*, 31(1):17–47, 1990.
- [8] T. Dieckmann and U. Schwalbe. Dynamic Coalition Formation and the Core, 1998. Economics Department Working Paper Series, Department of Economics, National University of Ireland - Maynooth.
- [9] E. Even-Dar, Y. Mansour, and U. Nadav. On the convergence of regret minimization dynamics in concave games. In *ACM Symposium on Theory of Computing (STOC)*, 2009.
- [10] D. Fotakis, S. Kontogiannis, and P. Spirakis. Atomic congestion games among coalitions. *ACM Trans. Algorithms*, 4(4):1–27, 2008.
- [11] X. Guo and H. Yang. The price of anarchy of Cournot oligopoly. In *Workshop on Internet and Network Economics (WINE)*, pages 246–257, 2005.

- [12] A. Harberger. Monopoly and resource allocation. *Amer. Econ. Review*, 44:77–87, 1954.
- [13] S. Hart and M. Kurz. Stable coalition structures. In M. J. Holler, editor, *Forming Coalitions*. Springer-Verlag, New York, 1982.
- [14] S. Hart and M. Kurz. Endogenous formation of coalitions. *Econometrica*, 51:1047–1064, 1983.
- [15] A. Hayrapetyan, É. Tardos, and T. Wexler. The effect of collusion in congestion games. In *ACM Symposium on Theory of Computing (STOC)*, pages 89–98, 2006.
- [16] J. Kluberg and G. Perakis. Generalized quantity competition for multiple products and loss of efficiency. In *46th Annual Allerton Conference on Communication, Control, and Computing*, pages 930 – 936, 2008.
- [17] L. Kóczy. A recursive core for partition function form games. *Theory and Decision*, 63:41–51, 2007.
- [18] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. *Computer Science Review*, 3(2):65–69, 2009.
- [19] A. Mas-Colell, J. Green, and M. D. Whinston. *Microeconomic Theory*. Oxford University Press, 1995.
- [20] U. Nadav and G. Piliouras. No regret learning in oligopolies: Cournot vs Bertrand. *3rd International Symposium on Algorithmic Game Theory (SAGT)*, 2010.
- [21] D. Ray. *A Game-Theoretic Perspective on Coalition Formation*. Oxford Univ. Press, 2007.
- [22] D. Ray and R. Vohra. Equilibrium binding agreements. *Journal of Economic Theory*, 73:30–78, 1997.
- [23] T. Roughgarden and E. Tardos. How bad is selfish routing? *J. ACM*, 49(2):236–259, 2002.
- [24] S. Suetens and J. Potters. Bertrand colludes more than Cournot. *Experimental Economics*, 10:71–77, 2007.
- [25] H. Young. *Strategic Learning and Its Limits*. Oxford University Press, Oxford, 2004.

Appendices

A Proof of Lemma 2

Here we present the proof of lemma 2.

Proof : Let's denote by Q the total production level, i.e. $Q = \sum_i^N q_i$. Similarly, for any coalition S_i , we define as $Q_{S_i} = \sum_{i \in S_i} q_i$. For each (super-)player S_i , we have that $u_{S_i}(\vec{q}) = \sum_{i \in S_i} (a - c - bQ)q_i = (a - c - bQ)Q_{S_i}$. In terms of social welfare we have that, $SW = \sum_{i=1}^k u_{S_i} = \sum_{i=1}^k (a - c - bQ)Q_{S_i} = (a - c - bQ)Q$. Similarly, in the case where, coalitions are replaced with single players with $q_i = Q_{S_i}$, we derive the same formula for the social welfare by the definition of Cournot games (i.e. $u_i(\vec{q}) = (a - c - bQ)q_i$).

If \vec{q} is a Nash equilibrium of the partition game, then for each i , $u_{S_i}(\vec{q}) = (a - c - bQ)Q_{S_i}$ is maximal given the choices of all the players not in S_i . Since, the utility of S_i depends only on the aggregate production level of the other coalitions, we have that Q_{S_i} maximizes $u_{S_i}(\vec{q})$, given the rest Q_{S_j} . Hence, $(Q_{S_1}, \dots, Q_{S_k})$ is a Nash equilibrium of a Cournot oligopoly with k players. Similarly, if $Q^* = (q_1, \dots, q_k)$ is a Nash equilibrium of a Cournot oligopoly with k players, then any vector $\vec{q} \in \mathcal{R}_+^{|S_i|}$ whose aggregate production levels for each coalition agree with their respective levels in Q^* is also a Nash equilibrium of the Cournot game with coalitions. Indeed, if that was not the case, there would exist a (super-)player with a deviating strategy such that her resulting utility would increase her utility given the strategies of the other (super-)players. Such a deviating strategy, must induce a differentiated aggregate production for (super-)player S_i , since all strategies with the same aggregate utility lead to the same payoff ($u_{S_i}(\vec{q}) = (a - c - bQ)Q_{S_i}$). Let's define this new aggregate production level as Q'_{S_i} , then in the original Cournot game Q'_{S_i} is a strictly better response for player i that her current strategy and we reach a contradiction. \square

B Proof of Lemma 3

Next, we present the proof of lemma 3.

Proof : From lemma 2, we have that the social welfare of all Nash equilibria of a Cournot oligopoly with a fixed coalition partition $\Pi = (S_1, \dots, S_k)$, is equal to the social welfare of the unique Nash of a Cournot oligopoly with k players. In the case of symmetric Cournot oligopolies, this social welfare is equal to $\frac{k}{(k+1)^2} \frac{(a-c)^2}{b}$. This is the case, because at the unique Nash the production level of all players is equal to $\frac{a-c}{(k+1)b}$. The optimal outcome arises when there is no competition (i.e. there exists a single producing player or coalition). In this case, the optimal payoff is equal to $\frac{1}{4} \frac{(a-c)^2}{b}$, since this is the global maximum of the function $\frac{(a-c-Q)Q}{b}$. By dividing the resulting social welfares we derive the price of anarchy. \square

C Proof of Theorem 11

Below, we present the proof of theorem 11.

Proof : By Lemma 3 we have that the cost of a partition with k coalitions is $O(k)$ times worse than that of optimal and that this ratio is increasing with k . Hence, the worst stable partition for any instance of a coalition formation game is the one with the maximal number of coalitions. Therefore, in order to bound the price of anarchy it suffices to find an upper bound on the number of coalitions k of a stable coalition structure. By Corollary 10, we thus need to optimize the following mathematical program:

$$\begin{aligned}
& \max && k \\
& \text{s.t.} && k = k_1 + k_2 \\
& && k_1 + \sum_{i=1}^{k_2} s_i = n \\
& && k_1 \leq (k_2 + 1)^2 \\
& && s_i \geq k^2 \quad \forall i \in \{1, \dots, k_2\} \\
& && k, k_1, k_2 \geq 0
\end{aligned}$$

We have that

$$n = k_1 + \sum_{i:s_i \geq 2} s_i \geq k_1 + k_2 \min_{i:s_i \geq 2} s_i \geq k_1 + k_2(k_1 + k_2)^2 \quad (1)$$

Trivially this implies that, $n \geq k_2^3$ or equivalently $k_2 \leq n^{\frac{1}{3}}$. Also, since by Corollary 10 we have that $k_2 \geq \sqrt{k_1} - 1$, equation (1) implies that $n \geq k_1 + (\sqrt{k_1} - 1)(k_1 + \sqrt{k_1} - 1)^2$. This is equivalent to $n \geq k^{\frac{5}{2}} + k^2 - 3k^{\frac{3}{2}} + 3k^{\frac{1}{2}} - 1$. As a result $k_1 \leq \max\{9, n^{\frac{2}{5}}\}$. Putting all these together we derive the desired bound on k and hence on the price of anarchy as well. \square

D Proof of Lemma 12

Here, we provide the proof of lemma 12. For any N , let n be the number: $n = \lceil 4N^{4/5} \rceil \lfloor N^{1/5} \rfloor + \lfloor N^{2/5} \rfloor$. For a game on n players, let Π be a partition of the n players into $k_1 = \lfloor N^{2/5} \rfloor$ singletons and $k_2 = \lfloor N^{1/5} \rfloor$ coalitions of size $s = \lceil 4N^{4/5} \rceil$ each. We will prove the desired Lemma, by exhaustively showing that for this coalition structure all types of deviating moves are not successful. We commence our analysis by looking into deviations of type 1.

Proposition 15 *There exists no successful deviation of Type 1 for Π .*

Proof : A deviation of type 1 is one in which a strict subset S'_i of a current coalition S_i decides to deviate and form a new coalition. The rest of the members of the original coalition (i.e. S_i/S'_i) dissolve into singletons. In our instance, the only such possible deviation is one where a number of $s' < s$ players deviate from a coalition of size s and form a new coalition with the remaining $s - s'$ members of the coalitions breaking off into singletons. The initial payoff of the players in that coalitions were $(a - c)^2 / (bs(k + 1)^2)$. After the deviation the payoffs for the deviating players will be equal to $(a - c)^2 / (bs'(k + (s - s') + 1)^2)$. In order for this coalition to not be successful it suffices that we have:

$$\begin{aligned}
(a-c)^2 / (bs(k+1)^2) &\geq (a-c)^2 / (bs'(k+(s-s')+1)^2) \Leftrightarrow \\
s'(k+1+(s-s'))^2 &\geq s(k+1)^2 \Leftrightarrow \\
s'((k+1)^2 + 2(k+1)(s-s') + (s-s')^2) &\geq s(k+1)^2 \Leftrightarrow \\
2(k+1)s'(s-s') + s'(s-s')^2 &\geq (s-s')(k+1)^2 \Leftrightarrow \\
2(k+1)s' + s'(s-s') &\geq (k+1)^2 \Leftrightarrow \\
-s'^2 + (2(k+1)+s)s' &\geq (k+1)^2 \Leftrightarrow
\end{aligned}$$

Given that $s' \in [1, s-1]$ the minimum of the quantity $X = -s'^2 + (2(k+1)+s)s'$ is achieved at either $s' = 1$ or $s' = s-1$. If $s' = 1$, then $X = s + 2k + 1 \geq k^2 + 2k + 1 = (k+1)^2$. Similarly, if $s' = s-1$, then $X = s'(-s' + 2(k+1) + s) = (s-1)(2k+3) \geq (k+1)^2$.

□

Similarly, we continue to argue that no deviations of type 2 are successful either.

Proposition 16 *There exists no successful deviation of Type 2 for Π .*

Proof : In our instance, the only such possible deviation is one where a number of $s' < s$ players deviate from a coalition and join either a singleton or another coalition of size s . In order to prove that such deviations are unsuccessful, we will argue that the members of the receiving coalition S_j (regardless of whether this is a singleton or a coalition of size s) are better off in the initial coalition structure. Indeed, in the new coalition structure, we will have $k + s - s' - 1 \geq k$ coalitions. Similarly, the size of the coalition S_j that receives the deviating players will strictly increase. Therefore, the resulting payoffs for the members of coalition S_j will decrease in any such deviation and as a result no such deviation can be successful. □

We conclude, by proving that deviations of type 3 cannot be successful. Deviations of type 3 allow for any number of coalitions to unite and form a new coalition. We exhaustively examine all possible compositions of the deviating sets and prove that there exists no set of coalitions for which a deviation of type 3 is profitable. Specifically, we have to examine the following cases for our instance:

- 1) All the coalitions that participate in the move are singletons.
- 2) All the coalitions that participate in the move are of size s .
- 3) Some of the participating coalitions are singletons and some are of size s .

It is straightforward to show that for the first two cases a deviation of type 3 cannot be successful. For the third case, we show that if any such deviation is successful, then there exists a successful deviation of type 3 in which all singletons participate. The proof concludes by showing that no such deviation can be successful either.

Proposition 17 *There exists no successful deviation of Type 3 for Π .*

Proof : We will show that no deviating moves of type 3 are successful, regardless of the composition of sets of coalitions participating in the move.

1) All the coalitions that participate in the move are singletons.

This essentially corresponds to a move of type 3 of the restricted coalition formation game of Section 3.1. By Lemma 8 and Lemma 9 it suffices to show that $k_1 \leq (k_2 + 1)^2$. This is true in our instance, since $\lfloor N^{2/5} \rfloor \leq N^{2/5} = (N^{1/5})^2 < (\lfloor N^{1/5} \rfloor + 1)^2$.

2) All the coalitions that participate in the move are of size s .

If we have m sets of size s uniting to form a single coalition, then the original payoff of the players were $(a - c)^2/s(k + 1)^2$, whereas their utilities in the new coalition will be equal to $(a - c)^2/ms(k - m + 2)^2$. Equivalently, if we denote $\lambda = k + 1$, the respective costs will be expressed as $(a - c)^2/s\lambda^2$ and $(a - c)^2/ms(\lambda - m + 1)^2$. In order to have stability it suffices that we have :

$$\begin{aligned}
(a - c)^2/s\lambda^2 &\geq (a - c)^2/ms(\lambda - m + 1)^2 \Leftrightarrow \\
ms(\lambda - m + 1)^2 &\geq s\lambda^2 \Leftrightarrow \\
m(\lambda^2 + m^2 + 1 - 2m\lambda + 2\lambda - 2m) &\geq \lambda^2 \Leftrightarrow \\
m\lambda^2 + m^3 + m - 2m^2\lambda + 2m\lambda - 2m^2 &\geq \lambda^2 \Leftrightarrow \\
(m - 1)\lambda^2 + (-2m^2 + 2m)\lambda + m^3 - 2m^2 + m &\geq 0 \Leftrightarrow \\
\lambda^2 - 2m\lambda + m(m - 1) &\geq 0 \Leftrightarrow \\
(\lambda - m - \sqrt{m})(\lambda - m + \sqrt{m}) &\geq 0 \Leftrightarrow
\end{aligned} \tag{2}$$

However, $\lambda = k + 1 = \lfloor N^{2/5} \rfloor + \lfloor N^{1/5} \rfloor + 1$, whereas $m \leq k_2 = \lfloor N^{1/5} \rfloor$. Hence, $\lambda \geq m + \sqrt{m}$, for all N .

3) Some of the participating coalitions are singletons and some are of size s .

We will start the analysis of this case by showing that if any such deviation is successful, then there exists a successful deviation of type 3 in which all the singletons participate.

Indeed, suppose that we have a profitable move of type 3 such that m_1 singletons and m_2 coalitions of size s participate, where $m_1 < k_1$. Each player, in the resulting coalition earns $(a - c)^2/b(m_1 + m_2s)(k - m_1 - m_2 + 2)^2$ and since this is a successful move, this payoff dominates the original payoff of the singleton players. Now, swap each coalitions of size s with singletons that do not participate currently in this move. If $m_1 + m_2 \leq k_1$, we will end up swapping all the non-singleton coalitions, so finally the sets that participate in the move are all singletons. As a result of this swapping the new payoffs for the players will be equal to $(a - c)^2/b(m_1 + m_2)(k - m_1 - m_2 + 2)^2$, which is strictly better than the profits when coalitions of size s were included. Hence, this move will be successful as well, but this is not possible, since this is essentially a move of type 3 of the restricted coalition formation game which has already been ruled out. The only case left, is the one where we have a move of type 3 with

m_1 singletons and m_2 coalitions of size s and $m_1 + m_2 > k_1$. As a result of the same swap, we define a move of type 3 with k_1 singletons and $m_1 + m_2 - k_1$ coalitions of size s . The payoff of the players in the resulting coalition will be $(a - c)^2 / (b(k + (m_1 + m_2 - k_1)s)(k - m_1 - m_2 + 2)^2)$, which dominates the payoffs of the first move of type 3 $(a - c)^2 / (b(m_1 + m_2)s)(k - m_1 - m_2 + 2)^2$. Hence, the new utilities of the singletons are higher than their utilities before the deviation and as result this deviation will be successful as well. So, if there exists a successful deviation of type 3, then there exists a successful deviation of type 3 where all the singletons participate (as well as some coalitions of size s).

At this point, in order to show that no moves of type 3 are successful for our family of examples, it suffices to show that there exist no successful moves of type 3 where we have all k_1 singletons participating as well $m \geq 1$ coalitions of size s . We will show that this move is not successful by showing that the singletons have no incentive to participate in this move.

Indeed, currently the singleton players enjoy a payoff of $(a - c)^2 / (b(k_1 + k_2 + 1)^2)$. After the deviation, their payoff will be equal to $(a - c)^2 / (b(k_1 + ms)(k_2 - m + 2)^2)$. In order for this move to not be successful, it suffices to have that $(k_1 + ms)(k_2 - m + 2)^2 \geq (k_1 + k_2 + 1)^2$.

$$\begin{aligned}
(k_1 + ms)(k_2 - m + 2)^2 &> 4s \\
&= 4\lceil 4N^{4/5} \rceil \\
&> (\lfloor N^{2/5} \rfloor + \lfloor N^{1/5} \rfloor + 1)^2 \\
&= (k_1 + k_2 + 1)^2
\end{aligned}$$

□

By combining propositions 15, 16 and 17, we have finally proven lemma 12.

E Proof of Theorem 14

Here, we present the proof of theorem 14 along with the relevant definitions. Hence, we will start by giving a formal definition of having no-regret in an online sequential problem.

Definition 7 *An online sequential problem consists of a feasible set $F \in \mathbb{R}^m$, and an infinite sequence of functions $\{f^1, f^2, \dots\}$, where $f^t : \mathbb{R}^m \rightarrow \mathbb{R}$.*

At each time step t , an online algorithm selects a vector $x^t \in \mathbb{R}^m$. After the vector is selected, the algorithm receives f^t , and collects a payoff of $f^t(x^t)$. All decisions must be made *online*, in the sense that an algorithm does not know f^t before selecting x^t , i.e., at each time t , a (possibly randomized) algorithm can be thought of as a mapping from a history of functions up to time t , f^1, \dots, f^{t-1} , to the set F .

Given an algorithm \mathcal{A} and an online sequential problem $(F, \{f^1, f^2, \dots\})$, if $\{x^1, x^2, \dots\}$ are the vectors selected by \mathcal{A} , then the payoff of \mathcal{A} until time T is $\sum_{t=1}^T f^t(x^t)$. The payoff of a static feasible vector $x \in F$, is $\sum_{t=1}^T f^t(x)$. Regret compares the performance of an algorithm with the best static action in hindsight:

Definition 8 *The external regret of algorithm \mathcal{A} , at time T is defined as*

$$\mathcal{R}(T) = \max_{x \in F} \sum_{t=1}^T f^t(x) - \sum_{t=1}^T f^t(x^t).$$

An algorithm is said to have no-external regret, if for every online sequential problem, its regret at time T is $o(T)$.

Next, we present the proof of theorem 14.

Proof : As we have argued in the proof of lemma 2, given any strategy profile of the Cournot game with coalitions the superplayers' payoff is exactly equal to the payoff of regular players that play the same game with production levels equal to the aggregate production of the corresponding superplayers.

If the superplayers experience no-regret, that means that deviating to any fixed strategy only decreases their average payoff¹. Specifically, since their payoff on any day depends only on the aggregate production, the average utility of any player only decreases when a player deviates and only plays strategies of a specific production level. Now, consider the same history of play but replace every superplayer by a simple player that produces on each day t quantity equal to the aggregate production quantity of the respective coalition on the same day. We will show that this history is also of no-regret. Indeed, for each player any deviation to a constant production level will only decrease their average utility¹, as is true for the superplayers. But, in the case of simple Cournot games, as was shown in [9], if all players experience no-regret then the average utilities of the players converge to their levels at Nash equilibrium. Translating this result to the history of the Cournot game with coalitions, we have that the average utilities of the superplayers converge to the levels of the unique Nash of the same Cournot game with k simple players. As was shown in lemma 2, these levels are also in agreement with the Nash equilibria of the Cournot game with coalitions. \square

¹More accurately, it means that at times goes to infinity the limsup of the difference between the average utility for any fixed deviation and the experienced average utility will be at most 0.