How to Play Out of Equilibrium: Beating the Play[*]

Karl H. Schlag

University of Vienna[§]

March 24, 2020

Abstract

Beating the play is a novel method for how to play a simultaneous move game without conjectures about how others play the game. A strategy beats the play if its payoff is higher than when playing like others play the game, regardless of how they play the game. Only Nash equilibrium strategies of the hypothetical game in which you play against copies of yourself can beat the play. It is possible to beat the play in numerous games. Many extensions are presented and a close connection to evolutionary game theory is uncovered.

Keywords: beating the play, solution concept, aggregative game, Cournot competition.

JEL classification number: C72.

1 Introduction

Nash equilibrium is the predominant solution concept for games and relies on the ability of each player to perfectly predict the behavior of the other players. However, often it is unknown what others do and there may be no common understanding of the underlying uncertainty.[1] Then we find ourselves in need of a strategy that

---

[*]I would like to thank Josef Hofbauer, Bernhard Kasberger, Simon Martin, Joel Sobel, Juha Tolvanen, Jörgen Weibull and seminar participants at University of Oxford for helpful comments.

[§]Department of Economics, University of Vienna, Oskar-Morgenstern-Platz 1, 1090 Vienna, Austria. *E-mail:* karl.schlag@univie.ac.at

[1]See also the Wilson (1987) critique that blames common knowledge for the disparity between theory and reality.

performs well when we do not expect that a Nash equilibrium or a Bayesian Nash equilibrium will be played. We approach the problem from an extreme position and propose a "personal" solution concept for a player who is entirely unable to predict what others will do. Later we expand the model to incorporate some understanding of, or conjectures about, the play of others. It is referred to as personal as we only make a recommendation for that player, not simultaneously for all players involved.

We develop a personal solution concept for how to play a game without making any assumptions how others play the game. In particular, there is no need to know the objectives or preferences of any of the other players. We outline the basic concept. We consider a player, who we refer to as player 0, who is about to play an n player simultaneous move game. This player perceives the different player roles as indistinguishable. So she could be facing any given way the game is being played from the perspective of any of the player roles. It is as if player 0 enters equally likely into any of the n player roles. It is as if she enters the game by replacing an existing player. With this scenario in mind, prior to the entry, consider the strategy profile that describes which strategy each of the n players wants to choose. In this profile, player 0 can evaluate her overall performance as an average of her performance in the different player roles. She could choose a best response to this average if she knew the strategy profile. However, she does not know the strategy profile and the same strategy cannot be a best response for each strategy profile unless player 0 has a weakly dominant strategy. So we lower the goal and ask the following question. Is there a single strategy for player 0 that is a better response for each strategy profile to the fictitious benchmark in which she chooses the strategy of the player she replaces? If yes, then we say that this strategy *beats the play*. Loosely speaking, this strategy performs better than when playing like others.

Note that the criterion underlying beating the play arises naturally when comparing performance of different strategies using data. Let's measure the performance of the strategy by considering its payoff against random opponents drawn from the data. Compare this to the performance of a random strategy from the data against random opponents from the data. To beat the play means that the given strategy does better than the random strategy in any data set.

We describe the concept of beating the play more formally. Let A be the set of pure

strategies of each of the n players and let u be the payoff (or utility) function of player 0. Let s_i denote the strategy chosen by the player that occupies the role of player i. To play like player i generates payoff $u(s_i, s_{-i})$ where $s_{-i} = (s_{i+1}, .., s_n, s_1, ...s_{i-1})$. It is as if players are located on a circle or that payoffs are anonymous in the sense that they do not depend on player indices. To play like others means to obtain payoff $\frac{1}{n} \sum_{i=1}^{n} u(s_i, s_{-i})$. To choose ξ as player i yields payoff $u(\xi, s_{-i})$. To choose ξ when facing the others means to obtain payoff $\frac{1}{n} \sum_{i=1}^{n} u(\xi, s_{-i})$. Then strategy ξ *beats the play* if $\frac{1}{n} \sum_{i=1}^{n} u(\xi, s_{-i}) \geq \frac{1}{n} \sum_{i=1}^{n} u(s_i, s_{-i})$ holds for all profiles $s \in A^n$. In particular, note that only the utility of player 0 matters, no information about that of the other players is needed.

A first insightful result is that a strategy that beats the play has to be a symmetric Nash equilibrium (NE) strategy in the so-called *game against selves*. This is the hypothetical symmetric game in which all players have the same payoff as you. The intuition is as follows. If all but one player use the same strategy as you do, then to perform better than when playing like others means to perform better than this odd man out when all others use your strategy. This is exactly the condition that ensures that all choosing your strategy is a NE of the game against selves. Thus, we find that it can make sense to choose a NE strategy even if you have no conjectures about what others will do. We also find that only pure strategies can beat the play under generic payoff functions.

We establish necessary and sufficient conditions for when a symmetric NE strategy of the game against selves beats the play when the set of actions A is a subset of $\mathbb{R}$. These build on the observation that a strategy beats the play if and only if $f_\xi(s) = \frac{1}{n} \sum_{i=1}^{n} u(\xi, s_{-i}) - \frac{1}{n} \sum_{i=1}^{n} u(s_i, s_{-i})$ is minimized when $s = \xi^n$. To obtain sufficient conditions we split this minimization up into first showing that f_ξ attains its smallest values on the diagonal, so where $s_i = s_j$ for i, j, and then showing that the minimum on the diagonal is attained when $s = \xi^n$. For the first argument we consider the minimization along the off diagonals where we keep the total sum of s_i constant. Easy to verify sufficient conditions are presented that involve convexity and concavity assumptions. Necessary conditions are formulated by applying the above minimization approach locally and lead to inequalities involving second derivatives. Under strategic substitutes we find that one only has to check the off diagonals. Under

strategic complements one only has to check the diagonal. In particular, neither strategic complements or substitutes are necessary or sufficient for beating the play.

We look into more detail at payoffs that only depend on own strategy and on the sum of the strategies as in an aggregative game and present many examples. Consider for instance Cournot competition. If own costs are convex, own demand is convex and own profits as monopolist are concave then the Cournot output beats the play. Note that payoffs exhibit strategic substitutes. Specifically, when there are n firms, you face inverse demand $\max\{1-q,0\}$ and have constant marginal costs c ($c < 1$) then set quantity $\frac{1-c}{n+1}$ to beat the play. Looking at data from laboratory experiments we discover that the Cournot output as the solution to beating the play actually realizes payoffs that are very close to those of the empirical best response. In the appendix (Section D.2) we explain this finding. Consider instead Bertrand competition with heterogenous goods in which own demand is convex and own costs exhibit constant marginal costs. Here too it is possible to beat the play and note that payoffs have strategic complements. Other examples in which one can beat the play include public good games and contests.

We see the property of beating the play as desirable but not as necessary. The concept itself is very demanding, not allowing to almost beat the play and requiring the property to hold for arbitrary strategy profiles. When it is not possible to beat the play then we offer generalizations of the basic definition that relax these aspects.

In our first generalization we drop the requirement that own payoff must be higher than it is when playing like the others. Instead we search for a strategy that never falls too far below playing like the others. Let the *shortcoming* of a strategy measure how far a strategy is from being able to beat the play. Specifically, the shortcoming of a strategy is defined as the maximal amount that the payoff from this strategy can lie below that obtained when playing like the others. We say that a strategy *comes closest to beating the play* if it attains the lowest shortcoming among all strategies. If the lowest shortcoming is small in a given context then we say that one can *almost beat the play*. We find that a strategy that comes closest to beating the play exists under minimal assumptions (for instance, when the set of actions A is a compact subset of $\mathbb{R}$ and utility u is continuous).

In some games the lowest shortcoming is small. This is for instance the case for

Bertrand competition with homogeneous goods and constant marginal costs. The strategy that comes closest to beating the play involves pricing above the competitive price with the extra margin decreasing in the number of firms. In particular, there is no Bertrand paradox. For instance, if there are n firms and you face linear demand $\max\{1-p,0\}$ and have constant marginal costs c then set price $\xi = c + \frac{1}{2}\left(1 - \sqrt{\frac{n}{n+1}}\right)(v-c)$ to come closest to beating the play. Its shortcoming is equal to $\frac{(1-c)^2}{4n(n+1)}$. Similarly, performance in first price auctions is very good. The strategy that comes closest to beating the play prescribes to bid the fraction $\frac{n}{n+1}$ of your own value v for the good. No assumptions on values or bidding behavior of others is made. Its shortcoming is equal to $\frac{v}{n(n+1)}$ which is arguably small when $n \geq 5$.

In other games one cannot almost beat the play as the shortcoming of any strategy is clearly large. For instance, this is the case when the game against selves is a coordination game. In such a case it is clear that one cannot perform well without some understanding about what others will be doing.

We proceed by providing a few extensions where we no longer check performance in all strategy profiles. First we adapt the definition so that a player never compares own performance to that of strictly dominated strategies.

Then we consider the setting where players choose their strategies independently. From the perspective of the player about to enter the game, it is as if strategies of opponents are independently drawn from the same unknown distribution. Accordingly, we say that ξ *beats the independent play* if $u\left(\xi,\sigma^{n-1}\right) \geq u\left(\sigma^n\right)$ holds for all mixed strategies σ. This condition equally emerges in the context of investigating data when ignoring correlation in play and matching strategies independently based on the empirical distribution. This definition can be shown to be equivalent to that of a globally neutrally stable strategy (Hofbauer & Sandholm, 2009) in the game against selves. We build on the results in this literature to find that one can beat the independent play in a war of attrition and in some congestion games.

In another variation we consider a hybrid model in which a player has some beliefs about which strategy profiles she is possibly facing. The condition to beat the play is only checked for strategy profiles (or distributions) belonging to this set U. We then say that ξ *beats the play in U*. This gives nice insights in settings where some

5

understanding of the play of others is needed. For instance, assume that the game against selves is a coordination game in which each symmetric pure strategy profile is a strict NE. Then any pure strategy beats the play whenever it is believed that others are sufficiently likely to use this strategy.

In the online appendix we also extend our definition to the case where player roles are not indistinguishable. Consider for instance a game between two types of players, such as buyers and sellers. A buyer will naturally only want to compare own performance to that of other buyers. One option is to drop sellers from the comparison, this is treated in Section A.1. An alternative pursued in Section A.2 is to introduce types which selects Bayesian Nash equilibria of the game against selves. Further extensions are provided to games where there is parametric uncertainty and where others also try to beat the play.

In this paper we uncover a close connection to evolutionary game theory, specifically to neutrally stable strategies. Under neutral stability, strategies are tested for whether they survive when competing with others. They are not actively chosen. Specifically, a strategy is called a *neutrally stable strategy* (NSS, Maynard Smith, 1982) if this strategy as incumbent does at least as well as a mutant strategy in a single infinite population where most use the incumbent strategy and only some use the mutant strategy. This is mathematically equivalent to our personal solution concept of beating the independent play in a neighborhood of this strategy. Dropping independence, our concept corresponds to a NSS of the symmetrized game introduced by Selten (1980) in which player roles are randomly assigned. When considering all strategy profiles, and not just those in a neighborhood, our concept is equivalent to that of a globally NSS (Hofbauer & Sandholm, 2009). Most of our examples are new to this literature. Our technique, to compute solutions through a minimization, and to do this separately on and off the diagonal, is novel. Our concept sheds a new light on neutral stability which currently is rarely used in economic applications as it is based on random matching in an infinite population which is typically very distant from economic environments that involve firms or workers. Despite the mathematical connections, the underlying motivations of neutral stability and beating the play are fundamentally different. Under neutral stability, different individuals are competing according to the payoffs they obtain. Under beating the play there is only a single

individual who is comparing different strategies.

In the literature only few solution concepts for games do not postulate that other players follow the same reasoning. Like us, some do not impose any restrictions on the play of others. Von Neumann and Morgenstern (1944) propose the security strategy (a maximin utility approach), it coincides with a NE strategy in a zero-sum game. Linhart and Radner (1989) look at minimax regret in Nash bargaining. Some restrictions are added by Kasberger and Schlag (2020) when they derive recommendations for bidding in first price auctions using a variant of minimax regret. Others investigate the outcomes when different types of reasoning are present among the players. Renou and Schlag (2011) investigate minimax regret when opponents follow the same reasoning with a given probability. Level k reasoning and its variants explore the interaction among players with different hierarchies of beliefs (Stahl, 1993, Camerer et al., 2004).

The most related concepts to beating the play are maximin utility and minimax regret (as well as any other concept for making decisions under ambiguity). As solution concepts for decision making they can be applied by treating the game as a decision problem. No assumptions about the payoffs in the roles of other players, unlike indistinguishability. However, without adding some conjectures about the play of others, maximin utility typically does not deliver sensible results (e.g., see Bertrand competition, Footnote 2.4 in Section 3.1.2). Minimax regret aims to find strategies that are always close to a best response to the play of others. It measures performance relative to an ambitious benchmark that makes it hard to even guarantee mediocre performance (for instance this is easily verified in Cournot competition[2]). On the other hand, beating the play evaluates performance relative to a modest benchmark and thus allows excellent performance in many games.

Finally, note that our replacement and comparison scenario is reminiscent of the boundedly rational approach in (Schlag, 1998). Therein, a single player replaces a random member of a population and learns from others to then make a choice in a decision problem. In the scenario where only strategies but not payoffs of others are observable then the objective therein can be formulated in terms of this player only.

---

[2]Interestingly, we find for Cournot competition in small markets that this maximal distance is small when computing performance in our entry scenario (see Section 2.5).

Payoffs of others then play no role.

We proceed as follows. In Section 2 we assume indistinguishable roles and explore beating the play. Therein, we motivate beating the play (Section 2.1), introduce the concept (Section 2.2), present some properties (Section 2.3), analyze aggregative games (Section 2.4) and investigate Cournot competition with kinks, including a data application (Section 2.5). Section 3 collects various extensions. In Section 3.1 we look at strategies that come closest to beating the play with an application to Bertrand competition (with a data application). In Section 3.2 we accommodate for the case where others use a strictly dominated strategy with an application to first price auctions. Then we limit the ways in which the game is being played, looking at independent play in Section 3.3 and play in a neighborhood in Section 3.4. In Section 4 we conclude. In the online appendix we expand our basic model in various directions. In Section A we show how to make recommendations when player roles are not indistinguishable. In Section A.1 the comparison is only to others that are indistinguishable from you. In Section A.2 we introduce types. In Section B we allow for parametric uncertainty, in Section C we consider multiple players simultaneously attempting to beat the play. In Section D we present two alternative definitions.

2   Beating the Play

Before we introduce our model in detail we provide a brief informal introduction to our personal solution concept, motivating it from the perspective of the reader.

2.1   Informal Motivation

Suppose you are about to play an n player simultaneous move game and need a recommendation for how to play this game. We present a novel methodology that does not require you to make any conjectures about what others do. Two key features define our innovation. The first identifies a new way to describe the strategic environment, the second introduces a new way to capture playing like others.

Assume that you cannot distinguish the different player roles. Later we will relax this condition. For a given way the game is being played you imagine you could equally

likely have faced this situation from the perspective of one of the other players. But if you think about being in the role of someone else then your seat becomes vacant and you cannot compute your payoffs when occupying a different role. We fill your seat with a strategy, it is as if each player has a strategy. Now you can analyze how your strategy performs from the perspective of each of the player roles. As you can imagine that you face the game from the perspective of each player role equally likely, it is as if you enter the game by randomly replacing one of the players. In fact, this is the first key to this paper. Instead of describing the environment you face as what strategies the $(n-1)$ other players choose we describe the environment you face as a profile of n strategies, one for each player.

The next step in the standard approach is to form beliefs about the environment and to best respond to these. We consider an alternative approach in which we aim to perform well in all environments. In particular we will design a benchmark and attempt to outperform this benchmark in each environment. The benchmark we introduce is the payoff you would achieve if you could choose the strategy of the player you replace. This is a purely hypothetical alternative. However it generates a good measure for the value of the strategies currently used. This is the second key of this paper, to define a benchmark that describes playing like others. Our objective is then to recommend a single strategy that performs better than when playing like others. So this approach applies if you do not want to rule out anything about how others play (in Section 3.4 we look at restricted play of others).

We provide some more motivating details by looking more specifically at a two player game. Let A be the set of pure strategies for each player and let u be your utility function. Assume that you are assigned to the role of player 1 and that it turns out that the other player, player 2, plays $a_2 \in A$. To play like player 1 means to play like the opponent of player 2. For you there is no difference between the roles of player 1 and player 2. In particular, you could have also found yourself in the role of player 2. Lets say that when you are in the role of player 2 then the other player, player 1, plays $a_1 \in A$. So to play like player 1 means to choose a_1 and to face a_2 and to get payoff $u(a_1, a_2)$. To play like player 2 means to choose a_2 and to face a_1 and to get payoff $u(a_2, a_1)$. As the two player roles are indistinguishable, it is as if you find yourself equally likely in the role of player 1 and of player 2. So to play in this

constellation like the player whose role you adapt means to get an expected payoff of $\frac{1}{2}u\left(a_1, a_2\right) + \frac{1}{2}u\left(a_2, a_1\right)$. This is the benchmark or counterfactual payoff to which we compare the expected payoff of the strategy we recommend. It is hypothetical as there is no formal entry or replacement taking place. In particular, this counterfactual payoff cannot be computed using observed data in a game in which you are taking part. It can be only computed if one uses data to evaluate expost how you would have performed in a game in which you did not take part in.

Now consider your performance if you choose the strategy ξ. When entering into role of player 1 you get $u\left(\xi, a_2\right)$, when in the role of player 2 then you get $u\left(\xi, a_1\right)$. So your expected payoff of choosing ξ in this constellation is given by $\frac{1}{2}u\left(\xi, a_2\right) + \frac{1}{2}u\left(\xi, a_1\right)$. To perform better or better respond to the player whose role you adapt means in the above constellation that $\frac{1}{2}u\left(\xi, a_2\right) + \frac{1}{2}u\left(\xi, a_1\right) \geq \frac{1}{2}u\left(a_1, a_2\right) + \frac{1}{2}u\left(a_2, a_1\right)$. To better respond regardless of how the game is played by others means that $\frac{1}{2}u\left(\xi, s_2\right) + \frac{1}{2}u\left(\xi, s_1\right) \geq \frac{1}{2}u\left(s_1, s_2\right) + \frac{1}{2}u\left(s_2, s_1\right)$ holds for all $s \in A^2$. If this is true we say that ξ *beats the play* and informally say that ξ performs better than when playing like others. Our recommendation to you is to choose a strategy that beats the play, provided such a strategy exists.

Three conditions are needed in order to apply this concept to a game. They are jointly motivated by the assumption that player 0 cannot distinguish the different player roles. Informally we identify these three conditions as player 0 regarding player roles as *indistinguishable*. First of all, the game has to be such that each player has the same set of pure strategies. If not, player roles could be told apart. Second of all, there needs to be a common method for how player 0 computes payoffs that can be applied to each player role. We call this *weak symmetry*. For example, it can be as if player 0 perceives players as located on a circle where the payoff in a player role is calculated with respect to the relative player positions of the others. Weak symmetry is needed to be able to claim that player 0 can see herself in any of the player positions. Third of all, when evaluating the performance of player 0 in a given environment as defined by a strategy profile, the different player roles are given equal weight. The equal weighting comes from the fact that no player role can receive a larger weight as otherwise it would be different from the rest.

2.2　The Definition

Let ΔB be the set of all distributions that have support in a set B. Let Γ be a normal form game with the following ingredients. There are n players. Each player has the same set of pure strategies, denoted by A. Letters i and j will be reserved for the player indices, so $i, j \in \{1, .., n\}$. Greek letter ξ (xsi) is reserved for our personal solution concept which is an element of ΔA. For $a \in A$ let $a^n \in A^n$ be such that $(a^n)_i = a$ for all $i \in \{1, .., n\}$. Let $s \in A^n$ be a typical profile of pure strategies where s_i is the pure strategy of player i. We also allow players to independently choose mixed strategies, a typical profile of mixed strategies is denoted by $\sigma \in (\Delta A)^n$. All examples in this paper are simultaneous move games, however the concepts apply equally to the normal form of any sequential game. Games in which different players have different action sets will be considered later.

We will not define payoffs of any of the players in this game. Throughout we will only be concerned with the payoffs of a player, who we refer to as player 0, who is going to play this game in one of the player roles. This means that any description of the characteristics of a game always refers only to how the game looks from the perspective of player 0. To simplify exposition, we will not mention each time that assumptions on payoffs in the game only refer to those of player 0. For example, if we write that demand is convex then we mean that the demand that player 0 faces is convex.

For $i \in \{1, ..., n\}$ let $u_{(i)}(s)$ be the payoff of player 0 when she is in the role of player i and the strategy profile is given by s, so player 0 is choosing s_i and player j is choosing s_j for $j \in \{1, ..., n\} \setminus \{i\}$. We assume that it is possible to translate how player 0 experiences the game in the role of player i into an equivalent experience when she has in the role of player 1. We adapt the term from Plan (2017).

Definition 1 *Player* 0 *regards player roles as* weakly symmetric *if for each* $i \in \{2, ..., n\}$ *that there exists a permutation* π_i *of* $\{1, ..., n\}$ *with* $\pi_i(1) = i$ *such that* $u_{(i)}(s) = u_{(1)}\left(s_i, s_{\pi_i(2)}..., s_{\pi_i(n)}\right).$

Unless mentioned otherwise we assume throughout that player 0 regards player roles as weakly symmetric. This does not necessarily mean that other players also

regard player roles as weakly symmetric as we make no assumptions on payoffs of others. We slightly abuse notation to simplify exposition and set $u \equiv u_{(1)}$ and $s_{-i}(s) = \left(s_{\pi_i(2)}..., s_{\pi_i(n)} \right)$. We then write $u(s_i, s_{-i})$ instead of $u_{(1)} \left(s_i, s_{\pi_i(2)}..., s_{\pi_i(n)} \right)$. Roles are weakly symmetric when players are located on a circle and payoffs are determined by how strategies are distributed around oneself. In that case $s_{-i} = (s_{i+1}, ..., s_n, s_1, ..., s_{i-1})$.

In the most common applications payoffs are anonymous in the sense that payoffs only depend on the own strategy and on those used by others, not on which of the others uses which strategy.

Definition 2 *Player* 0 *regards player roles as* (strategically) anonymous *if* $u_{(i)}(s) = u_{(i)} \left(s_{\pi(1)}, s_{\pi(2)}..., s_{\pi(n)} \right)$ *holds for any permutation* π *of* $\{1, ..., n\}$ *with* $\pi(i) = i$ *for each* $i \in \{1, ..., n\}$.

Clearly, anonymity is stronger than weak symmetry. Both concepts coincide when $n = 2$.

As we can map payoffs in any role to payoffs that would be obtained in the role of player 1 we can define dominance independent of the player role. Accordingly, the strategy $a \in A$ is *strictly dominated* if there exists $x \in \Delta A$ such that $u(x, s_{-1}) > u(a, s_{-1})$ for all $s \in A^n$. On the other hand, $x \in \Delta A$ is a *weakly dominant* strategy if $u(x, s_{-1}) \geq u(s_i, s_{-i})$ for all $s \in A^n$.

Let Γ^0 be the hypothetical normal form game in which each player chooses a pure strategy from A and where player i has utility $u_i(s) = u(s_i, s_{-i})$, $i = 1, ..., n$. We refer to Γ^0 as the *game against selves*. We call $a \in \Delta A$ a *symmetric NE strategy* of Γ^0 if $a^n = (a, ..., a)$ is a NE of Γ^0.[3]

We assume that player 0 describes the environment she faces by a strategy profile $s \in A^n$ and evaluates her use of the strategy $x \in \Delta A$ in this environment by the expected payoff she achieves when randomly assigned (with equal probability) a player role i and facing s_{-i}. Formally, this payoff is given by $\frac{1}{n} \sum_{i=1}^n u(x, s_{-i})$. Our concept

---

[3]As player roles are assumed indistinguishable, it makes sense to refer to Γ^0 as a symmetric game. This is consistent with the literature when $n = 2$. However there are different definitions of symmetry in the literature when $n \geq 3$. Following Plan (2017), Γ^0 is weakly symmetric if $n \geq 3$ and a totally symmetric if player roles are also anonymous.

is based on comparing this payoff to the expected payoff she would hypothetically achieve if she could choose the strategy s_i of the player she replaces, as identified by $\frac{1}{n} \sum_{i=1}^{n} u(s_i, s_{-i})$. This leads us to the central concept of this paper.

Definition 3 $\xi \in \Delta A$ beats the play *if*

$$f_\xi(s) := \frac{1}{n} \sum_{i=1}^{n} \left(u(\xi, s_{-i}) - u(s_i, s_{-i}) \right) \geq 0 \ holds \ for \ all \ s \in A^n.$$

Performance is measured by considering a given strategy profile from the perspective of each player and averaging over the different roles. It is as if player 0 replaces equally likely one of the existing players and evaluates her strategy from the perspective of the player she replaced. An interpretation in terms of doing better than the average payoff when all players have identical preferences is misleading. For one thing, it is unrealistic to assume that players have identical preferences. Assuming nevertheless identical preferences opens the door for eliminating strategy profiles as plausible which is not possible when payoffs of others are not known. Our focus is on the setting where the payoffs of others are not known (or ignored), which consequently leads to a definition that only concerns the payoffs of the player of interest.

Remark 1 *Note that Definition 3 would not have changed if we had included beliefs over different profiles and demanded that $\int f_\xi(s) \, dG(s) \geq 0$ holds for all distributions G over A^n.*

Endowed with a strategy ξ that beats the play one does not have to worry about how the game is being played. The performance of ξ is guaranteed by the mathematics and does not require an evaluation based on experience. However, one might never the less wish to calculate $f_\xi(s)$ given some data. We comment. For a player who is actually playing the game, say as player 1, the entire profile s is never observable. The value of s_1 is only a thought construct used to evaluate performance when in the shoes of others. One could however compute $f_\xi(s)$ as a function of s_1. For an expost analysis based on a given data set to which the player of interest did not contribute, the distribution of profiles is observable, which makes $f_\xi(s)$ computable.

The following connection to evolutionary game theory is noteworthy. It will be used to provide some independent insights to those working on evolutionary game

theory but is not useful to derive or understand any of our main arguments. Hence we only briefly mention this connection. Accordingly, ξ beats the play if and only if ξ^n is a globally neutrally stable strategy in the symmetrized game of the game against selves. So first one considers the game against selves in which all players have payoff u. Then one symmetrizes this game (as in Selten, 1980) by randomly allocating each player to a role in this game. Finally one computes a globally neutrally stable strategy (Hofbauer & Sandholm, 2009) of this symmetrized game.

In the next sections we will investigate beating the play in games. We hasten to point out that this concept is restrictive in at least three different dimensions. There is no tolerance for performing slightly worse than when playing like others. There is no room for asymmetries as all player roles are indistinguishable. There are no beliefs as the condition is required for all strategy profiles. Later we relax the definition to address these different aspects. In the appendix we briefly discuss an alternative definition in which the objective is to beat the worst strategy being played (see Section D.1). In the appendix we also present a methodology for calculating the magnitude by which a strategy beats the play (see Section D.2).

2.3   Understanding the Concept

Clearly, strategies that beat the play need not exist. An example is the pure coordination context which is explicitly analyzed in Section 3.4. It is more astonishing that the concept exists in so many games as shown in Section 2.4. However, before getting to these examples we present some general findings.

We relate our concept to rationality and NE.

Proposition 1 *Assume that ξ beats the play. Then ξ is a symmetric NE strategy of Γ^0, in particular, ξ is not strictly dominated.*

Proof. Assume that $u\left(z, \xi^{n-1}\right) > u\left(\xi^n\right)$ for some $z \in A$. Then

$$f_\xi\left(z, \xi^{n-1}\right) = \frac{1}{n}\left(u\left(\xi^n\right) - u\left(z, \xi^{n-1}\right)\right) < 0$$

and hence, following Remark 1, ξ does not beat the play. ∎

In the context of part (i) above we hasten to point out the following. The defining properties of a strategy that beats the play does not involve that others use this strategy. This stands in contrast to the definition of a symmetric NE strategy that is only applicable as a prediction when all others use this strategy. In Section C we extend our concept to the case where it is known that others are also attempting to beat the play.

The connection to NE comes at no surprise to a reader who is familiar with evolutionary game theory. Recall from above that a strategy that beats the play is associated to a neutrally stable strategy of a symmetrized game. As neutrally stable strategies are NE strategies the result follows. Our next result can similarly connected to this literature. Bhaskar (1995) shows for two player games that neutrally stable strategies of the symmetrized game are generically pure strategies. We present a different proof for $n \geq 2$ players. Our proof nicely reveals the particular type of genericity under which only pure strategies can beat the play.

Proposition 2 *Let A be finite. Let player roles be anonymous and let ξ be a symmetric NE strategy of the game against selves. If $u\left(\xi^{k}, a^{n-k}\right) \neq u\left(\xi^{n}\right)$ for some $a \in C\left(\xi\right)$ and $k \in \{1, ..., n-1\}$ and ξ beats the play then ξ is a pure strategy.*

Proof. Assume that ξ beats the play and that ξ is not a pure strategy. Consider $a \in C\left(\xi\right)$. So $a \neq \xi$. Let $\mu \in \left(0, \xi\left(a\right)\right)$ and $b \in \Delta A$ be such that $\xi = \left(1 - \mu\right)\left[b\right] + \mu\left[a\right]$. Consider the game $\bar{\Gamma}$ with action set $\{a, b\}$ and payoff function $\bar{u}\left(s\right) = u\left(\left(\left(1 - \mu\right)\left[\xi\right] + \mu\left[s_{i}\right]\right)_{i=1}^{n}\right)$ for $s \in \{a, b\}^{n}$. Note that ξ beats the play in $\bar{\Gamma}$ under $\bar{u}$.

As ξ beats the play, $\bar{u}\left(\xi, a^{n-1}\right) \geq \bar{u}\left(a, a^{n-1}\right)$ and hence $\bar{u}\left(b, a^{n-1}\right) \geq \bar{u}\left(a, a^{n-1}\right)$.

Assume $\bar{u}\left(b, b^{k-1}, a^{n-k}\right) > \bar{u}\left(a, b^{k-1}, a^{n-k}\right)$ for some $k \in \{1, .., n-2\}$. So $\bar{u}\left(b, b^{k-1}, a^{n-k}\right) > \bar{u}\left(\xi, b^{k-1}, a^{n-k}\right)$. Let us show that $\bar{u}\left(b, b^{k}, a^{n-k-1}\right) > \bar{u}\left(a, b^{k}, a^{n-k-1}\right)$. Adding a superscript to f to identify the context we obtain

$$
\begin{aligned}
n f_{\xi}^{\{a,b\}, \bar{u}}\left(b^{k}, a^{n-k}\right) \;=\; & k\left(\bar{u}\left(\xi, b^{k-1}, a^{n-k}\right) - \bar{u}\left(b, b^{k-1}, a^{n-k}\right)\right) \\
& + (n-k)\left(\bar{u}\left(\xi, b^{k}, a^{n-k-1}\right) - \bar{u}\left(a, b^{k}, a^{n-k-1}\right)\right).
\end{aligned}
$$

Above we showed that the first term is strictly negative. As ξ beats the play in $\bar{\Gamma}$ under $\bar{u}$ it follows that $\bar{u}\left(\xi, b^{k}, a^{n-k-1}\right) > \bar{u}\left(a, b^{k}, a^{n-k-1}\right)$. This means that

15

$\bar{u}\left(b, b^k, a^{n-k-1}\right) > \bar{u}\left(a, b^k, a^{n-k-1}\right)$ which shows our claim. Repeating this argument shows that $\bar{u}\left(b, b^{n-1}\right) > \bar{u}\left(a, b^{n-1}\right)$. This then implies $\bar{u}\left(b, b^{n-1}\right) > \bar{u}\left(\xi, b^{n-1}\right)$ which is a contradiction to ξ beating the play. Consequently we have shown that $\bar{u}\left(b, b^{k-1}, a^{n-k}\right) = \bar{u}\left(a, b^{k-1}, a^{n-k}\right)$ holds for all $k \in \{1, .., n\}$. In particular, $\bar{u}\left(\xi, a^{n-1}\right) = \bar{u}\left(\xi^n\right) = u\left(\xi^n\right)$.

So we have shown that $\bar{u}\left(\xi, a^{n-1}\right) = u\left(\xi, ((1-\mu)[\xi] + \mu[a])^{n-1}\right) = u\left(\xi^n\right)$ holds for all $\mu \in (0, \xi(a))$. Using the identity theorem for polynomials we obtain $u\left(\xi^{n-k}, a^k\right) = u\left(\xi^n\right)$ for $1 \leq k \leq n-1$. This is a contradiction to the genericity assumption and hence completes the proof. ∎

We provide an example of a game in which a mixing is required in order to beat the play. Consider payoffs in the game against selves as in the Rock-Scissor-Paper game with $w = l$ (see Example 2.3 in Hofbauer & Sandholm, 2009).[4] To beat the play in this game requires to mix with equal weight between the three pure strategies. This follows from the observation that a symmetric NE strategy of the game against selves beats the play if $n = 2$ and $u(a, a') + u(a', a) = 0$ holds for all $a, a' \in A$.[5]

Next we compile useful sufficient conditions for a pure strategy to beat the play. We use the fact that ξ beats the play if and only if ξ^n minimizes $f_\xi(s)$ over all $s \in A^n$. The idea is to verify this condition in two steps. First we show that f_ξ attains its minimum on the diagonal $\{s \in A^n : s_i = s_j \text{ for all } i \neq j\}$. Then we show that f_ξ is lowest on the diagonal at ξ. When we consider whether the minimum is on the diagonal we move in a direction that keeps the sum of the strategies constant. Under appropriate convexity conditions one does not have to additionally verify that the minimum is on the diagonal as this follows from anonymity. Similarly, convexity on the diagonal together with ξ being a symmetric NE strategy of Γ^0 is sufficient to show that the minimum is attained at ξ^n.

---

[4]$A = \{R, S, P\}$, $u(a, a) = 0$ for $a \in A$, $u(a, a') = w$ if $(a, a') \in \{(R, S), (S, P), (P, R)\}$ and $u(a, a') = l$ otherwise.

[5]We verify

$$\frac{1}{2}\left(u(\xi, a) + u(\xi, b) - u(a, b) - u(b, a)\right) = \frac{1}{2}\left(u(\xi, a) + u(\xi, b)\right) = -\frac{1}{2}\left(u(a, \xi) + u(b, \xi)\right) \geq -u(\xi, \xi) = 0$$

Proposition 3 *Assume $A \subseteq \mathbb{R}$ and A is convex.*

(i) ξ beats the play if

$$u\left(\xi, a^{n-1}\right) \geq u\left(a^n\right) \text{ for all } a \in A, \text{ and} \tag{1}$$

$$f_\xi\left(s\right) \geq f_\xi\left(\left(\frac{1}{n}\sum_{i=1}^{n} s_i\right)^n\right) \text{ for all } s \in A^n. \tag{2}$$

(ii) Assume that player roles are anonymous. Then $\xi \in A$ beats the play if ξ is a symmetric NE strategy of Γ^0, u is differentiable at ξ^n,

$$u\left(\xi, a^{n-1}\right) - u\left(a^n\right) \text{ is convex in } a \text{ for all } a \in A, \text{ and} \tag{3}$$

$$f_\xi\left(s + \lambda\left(e_i - e_j\right)\right) \text{ is convex in } \lambda \text{ in a neighborhood of } \lambda = 0 \text{ for all } s \in A^n. ^{6}\tag{4}$$

If f_ξ is twice differentiable then (4) is equivalent to

$$\frac{\partial^2}{\left(\partial s_i\right)^2} f_\xi\left(s\right) - 2\frac{\partial}{\partial s_i}\frac{\partial}{\partial s_j} f_\xi\left(s\right) + \frac{\partial^2}{\left(\partial s_j\right)^2} f_\xi\left(s\right) \geq 0 \text{ for all } s \in A^n. \tag{5}$$

Note that a sufficient condition for (3) is that $u\left(\xi, a^{n-1}\right)$ is convex in a and $u\left(a^n\right)$ is concave in a.[7]

Proof. Part (i). (2) implies that f_ξ attains its minimum on the diagonal. (1) implies that the minimum on the diagonal is attained at ξ^n.

Part (ii). We apply part (i).

We first show that the minimum of f_ξ on the diagonal is attained at ξ^n. Let $g\left(a\right) = f_\xi\left(a^n\right)$. So we aim to show that $g\left(a\right) \geq g\left(\xi\right)$ for all $a \in A$. To prove this claim, note that $g\left(a\right) = u\left(\xi, a^{n-1}\right) - u\left(a^n\right)$. Following (3), g is convex in a. Moreover, $g'\left(\xi\right) = -\frac{\partial}{\partial s_1}u\left(s\right)|_{s=\xi^n}$ as u_i is differentiable at ξ^n. Assume that ξ belongs to the interior of A. Then as ξ is a symmetric NE strategy of Γ^0 and u is differentiable at ξ^n, $\frac{\partial}{\partial s_1}u\left(s\right)|_{s=\xi^n} = 0$. Hence, $g'\left(\xi\right) = 0$. Together with the convexity of g the claim follows. Assume now that ξ is at the left boundary of A. Then $\frac{\partial}{\partial s_i}u_i\left(\xi^n\right) \leq 0$ and hence $g'\left(\xi\right) \geq 0$. Together with the convexity of g this proves our claim. Similarly the statement follows when ξ is on the right boundary of A.

We now show that f_ξ attains its minimum on the diagonal. Fix $s \in A^n$. Consider $i \neq j$ with $s_i \neq s_j$. Let $g\left(x\right) = f_\xi\left(s + x\left(e_i - e_j\right)\right)$ for $x \in M = \left\{x : s + x\left(e_i - e_j\right) \in A^n\right\}$. Note that M is convex and contains $\frac{s_j - s_i}{2}$ in its interior. (4) implies that g is convex.

---

[7]Note that if u is twice differentiable then $\frac{d}{da}\frac{d}{da}u\left(\xi, a^{n-1}\right) = \left(n-1\right)^2\frac{\partial}{\partial s_2}\frac{\partial}{\partial s_2}u\left(\xi, a^{n-1}\right)$.

Note that $\left(s + \frac{s_j - s_i}{2}\left(e_i - e_j\right)\right)_i = \left(s + \frac{s_j - s_i}{2}\left(e_i - e_j\right)\right)_j$. Symmetry of u implies that $g\left(\frac{s_j - s_i}{2} + \varepsilon\right) = g\left(\frac{s_j - s_i}{2} - \varepsilon\right)$ when $\left\{\frac{s_j - s_i}{2} + \varepsilon, \frac{s_j - s_i}{2} - \varepsilon\right\} \subset M$. Consequently, g attains its minimum at $\frac{s_j - s_i}{2}$. As equalization of the i-th and j-th component in s strictly reduces the variance of s it follows that f_ξ attains its minimum on the diagonal.

Finally, (5) is easily verified. ∎

Finally, we present necessary conditions for a pure strategy to beat the play. These are derived when moving on and off the diagonal at ξ^n, yielding local versions of (3) and (4). First note that the first and second order necessary conditions for ξ to be a symmetric NE strategy of Γ^0 are $\frac{\partial}{\partial s_1} u\left(\xi^n\right) = 0$ and $\frac{\partial^2}{(\partial s_1)^2} u\left(\xi^n\right) \leq 0$.

Proposition 4 *Assume that A is a convex subset of $\mathbb{R}$ and player roles are anonymous. If $\xi \in int\left(A\right)$ beats the play and u is twice continuously differentiable in a neighborhood of ξ^n then*

$$\frac{\partial^2}{(\partial s_1)^2} u\left(\xi^n\right) \leq 2\frac{\partial}{\partial s_1}\frac{\partial}{\partial s_2} u\left(\xi^n\right) \leq -\frac{1}{n-1}\frac{\partial^2}{(\partial s_1)^2} u\left(\xi^n\right).$$

Here is some intuition, the formal proof is below. Beating the play refers to differences. It is as if one starts with profile s near ξ^n and one player deviates to ξ. Marginally this difference is equal to 0 due to first order condition for ξ^n to be a NE. Close to ξ^n we will be changing these differences marginally, hence the necessary condition refers to second derivatives.

Consider a change on the diagonal where current play moves from ξ^n to x^n with $x > \xi$. The direct effect is driven by the player you are replacing choosing a higher strategy, which is as if you are choosing a lower strategy. The marginal difference is $-\frac{\partial}{\partial s_1} u\left(x^n\right)$, the marginal change equals $-\frac{\partial}{\partial s_1}\frac{\partial}{\partial s_1} u\left(x^n\right)$. This is positive at ξ^n by the second order condition associated to ξ^n being a NE. The indirect effect is driven by change in behavior of the others. Here cross derivatives matter. Under strategic substitutes (where $-\frac{\partial}{\partial s_1}\frac{\partial}{\partial s_2} u\left(\xi^n\right) \geq 0$) the increase in their strategy increases your performance difference. Under strategic complements (where $-\frac{\partial}{\partial s_1}\frac{\partial}{\partial s_2} u\left(\xi^n\right) \leq 0$) the difference is decreased and needs to be watched. The magnitude of the decrease is $\frac{\partial}{\partial s_1}\frac{\partial}{\partial s_2} u\left(\xi^n\right)$. This countervailing effect has to be multiplied by $(n-1)$ as it is driven separately by each of the opponents and has to be doubled as it comes with a change both in your performance as well as in that of the benchmark of playing like others.

So the positive effect driven by concavity in own strategy has to be stronger than $2(n-1)$ times the negative effect driven by strategic complements. This leads to the necessary inequalities

$$-\frac{\partial^2}{(\partial s_1)^2}u\left(\xi^n\right) \geq 2\left(n-1\right)\frac{\partial}{\partial s_1}\frac{\partial}{\partial s_2}u\left(\xi^n\right).$$

Now consider a change on the off diagonal. Starting at ξ^n, assume that player j moves to $\xi + \varepsilon$, player $k \neq j$ moves to $\xi - \varepsilon$ and the rest remain at ξ. Say you replace player j. Again, the direct effect on the difference, driven by the different behavior of you and the player who you replace, is given by $-\frac{\partial^2}{(\partial s_1)^2}u\left(\xi^n\right) \geq 0$. The indirect effect is driven by player k changing her strategy in the opposite direction, moving to $\xi - \varepsilon$. This means that the performance difference is increased under strategic complements and decreased on strategic substitutes. As on the diagonal, the change has to be doubled. Consequently, concavity in own strategy has to be stronger than twice the negative effect driven by strategic substitutes. This leads to the necessary inequalities

$$-\frac{\partial^2}{(\partial s_1)^2}u\left(\xi^n\right) \geq -2\frac{\partial}{\partial s_1}\frac{\partial}{\partial s_2}u\left(\xi^n\right).$$

In short, on the diagonal the forces go in opposite direction which might pose a problem under strategic complements. On the off diagonal the opposite happens, forces go in the same direction which might cause a problem under strategic substitutes. Consequently, for strategic complements (so $\frac{d}{ds_1}\frac{d}{ds_j}u \geq 0$) we only have to check the diagonal, while for strategic substitutes (so $\frac{d}{ds_1}\frac{d}{ds_j}u \leq 0$) we only have to check the off diagonal. In particular, neither strategic complements nor strategic substitutes are necessary for beating the play.

Proof. We prove the first inequality that refers to the off diagonal by looking at changes in two entries of ξ^n that keep the sum of strategies constant. Fix $j, k \in \{1, ..., n\}$ with $j \neq k$. Let $h_i\left(s\right) = u\left(\xi, s_{-i}\right) - u\left(s_i, s_{-i}\right)$, $g_i\left(x\right) = h_i\left(\xi + x\left(e_k - e_j\right)\right)$ and $g\left(x\right) = nf_\xi\left(\xi + x\left(e_k - e_j\right)\right) = \sum_{i=1}^n g_i\left(x\right)$ for x in a neighborhood of 0 and $s \in A^n$. As ξ beats the play, $x = 0$ is a minimum of g so we need $g''\left(0\right) \geq 0$.

We compute $g''\left(0\right)$. Note that $g_i''\left(0\right) = 0$ if $i \notin \{j, k\}$. For $i \in \{j, k\}$ we have

$$g_i'\left(0\right) = \frac{\partial}{\partial s_k}h_i\left(\xi^n\right) - \frac{\partial}{\partial s_j}h_i\left(\xi^n\right),$$

$$g_i''\left(0\right) = \frac{\partial}{\partial s_k}\frac{\partial}{\partial s_k}h_i\left(\xi^n\right) + \frac{\partial}{\partial s_j}\frac{\partial}{\partial s_j}h_i\left(\xi^n\right) - 2\frac{\partial}{\partial s_j}\frac{\partial}{\partial s_k}h_i\left(\xi^n\right).$$

19

where

$$\frac{\partial}{\partial s_j} h_j(s) = -\frac{\partial}{\partial s_j} u(s_j, s_{-j}),$$

$$\frac{\partial}{\partial s_j}\frac{\partial}{\partial s_j} h_j(s) = -\frac{\partial}{\partial s_j}\frac{\partial}{\partial s_j} u(s_j, s_{-j}),$$

$$\frac{\partial}{\partial s_j}\frac{\partial}{\partial s_k} h_j(s) = -\frac{\partial}{\partial s_j}\frac{\partial}{\partial s_k} u(s_j, s_{-j}),$$

$$\frac{\partial}{\partial s_k} h_j(s) = \frac{\partial}{\partial s_k} u(\xi, s_{-j}) - \frac{\partial}{\partial s_k} u(s_j, s_{-j}),$$

$$\frac{\partial}{\partial s_k}\frac{\partial}{\partial s_k} h_j(s) = \frac{\partial}{\partial s_k}\frac{\partial}{\partial s_k} u(\xi, s_{-j}) - \frac{\partial}{\partial s_k}\frac{\partial}{\partial s_k} u(s_j, s_{-j}).$$

So $g''(0) = g_j''(0) + g_k''(0)$ where

$$g_j''(0) = -\frac{\partial}{\partial s_j}\frac{\partial}{\partial s_j} u_j(\xi^n) + 2\frac{\partial}{\partial s_j}\frac{\partial}{\partial s_k} u_j(\xi^n),$$

and hence $g''(0) \geq 0$ implies

$$-\frac{\partial}{\partial s_i}\frac{\partial}{\partial s_i} u_i(\xi^n) + 2\frac{\partial}{\partial s_i}\frac{\partial}{\partial s_j} u_i(\xi^n) \geq 0 \text{ for } j \neq i.$$

We prove the second inequality that refers to the diagonal by looking at changes near ξ on the diagonal. Let $g(x) = f_\xi(x^n) = u(\xi, x^{n-1}) - u(x^n)$ for $x \in A$. Then we need $g''(\xi) \geq 0$. We find

$$g'(x) = -\frac{\partial}{\partial s_1} u(x^n) + \sum_{j \neq 1}\left(\frac{\partial}{\partial s_j} u(\xi, x^{n-1}) - \frac{\partial}{\partial s_j} u(x^n)\right),$$

$$g''(\xi) = -\sum_k \frac{\partial}{\partial s_k}\frac{\partial}{\partial s_1} u(\xi^n) + \sum_{j \neq 1}\frac{\partial}{\partial s_1}\left(-\frac{\partial}{\partial s_j} u(\xi^n)\right)$$

$$= -\frac{\partial}{\partial s_1}\frac{\partial}{\partial s_1} u(\xi^n) - 2(n-1)\frac{\partial}{\partial s_1}\frac{\partial}{\partial s_2} u(\xi^n),$$

and hence $g''(0) \geq 0$ implies

$$-\frac{\partial}{\partial s_1}\frac{\partial}{\partial s_1} u(\xi^n) - 2(n-1)\frac{\partial}{\partial s_1}\frac{\partial}{\partial s_2} u(\xi^n) \geq 0.$$

∎

2.4   Differentiable Aggregative Payoffs

Consider now games in which player roles are anonymous and payoffs of player 0 are as in an aggregative game (Acemoglu & Jensen, 2013). Assume additionally that

player roles are anonymous and payoffs are twice differentiable. So A is a convex subset of $\mathbb{R}$ and there is some twice differentiable function $F : \mathbb{R}^2 \to \mathbb{R}$ such that

$$u(s) = F\left(s_1, \sum_{j=1}^{n} b(s_j)\right)$$

where b is strictly increasing.

Note that there is an equivalent representation in which the sum does not depend on $j = 1$, so $u(s) = F_0\left(s_1, \sum_{j=2}^{n} b(s_j)\right)$ for some F_0. As b is invertible we can interpret $b(a)$ instead of a as the action. This leads to the equivalent representation we will be using in which

$$u(s) = U(s_1, \Sigma(s))$$

for some U where U is twice differentiable and $\Sigma(s) = \sum_{j=1}^{n} s_j$.

Note that strategies are *strategic complements* if $U_{12}(x, y) + U_{22}(x, y) \geq 0$, and *strategic substitutes* if $U_{12}(x, y) + U_{22}(x, y) \leq 0$ for all x, y, where $U_{kl}(x_1, x_2) = \frac{\partial}{\partial x_k}\frac{\partial}{\partial x_l}U(x_1, x_2)$.[8]

The necessary conditions for ξ to be a symmetric NE strategy of Γ^0 are that $U_1(\xi, n\xi) + U_2(\xi, n\xi) = 0$ and $U_{11}(\xi, n\xi) + 2U_{12}(\xi, n\xi) + U_{22}(\xi, n\xi) \leq 0$. The necessary conditions presented in Proposition 4 for $\xi \in A$ to beat the play turn into

$$
\begin{aligned}
U_{11}(\xi, n\xi) &\leq U_{22}(\xi, n\xi), \text{ and} && (6)\\
U_{12}(\xi, n\xi) + U_{22}(\xi, n\xi) &\leq \frac{1}{2n}(U_{22}(\xi, n\xi) - U_{11}(\xi, n\xi)).
\end{aligned}
$$

The first inequality refers to conditions derived on the off diagonal, the second to those derived on the diagonal.

We use Proposition 3 to derive sufficient conditions for a pure strategy to beat the play.

Proposition 5 *Let $\xi \in A$ be a symmetric NE strategy of Γ^0. Then ξ beats the play if*

$$
\begin{aligned}
(n-1)^2 U_{22}(\xi, \xi + (n-1)x) &\geq \frac{d}{dx}\frac{d}{dx}U(x, nx), && (7)\\
U_{22}(\xi, \xi + \Sigma(s_{-1})) &\geq U_{11}(s_1, \Sigma(s)). && (8)
\end{aligned}
$$

---

[8]More precisely, player 0 perceives the strategies like this. However, following the general statement at the beginning of Section 2.2 we refrain from mentioning that this is actually only about player 0.

(7) and (8) hold if $U(x,y)$ is concave in x, $U(\xi,y)$ is convex in y and $U(x,nx)$ is concave in x.[9]

Proof. We apply Proposition 3 (ii). Let

$$h(x) = U(\xi, \xi + (n-1)x) - U(x, nx).$$

Then (7) implies that h is convex and hence that (3) holds.

We now verify (4). Let

$$h_i(s) := u(\xi, s_{-i}) - u(s_i, s_{-i}) = U(\xi, \xi + \Sigma(s_{-i})) - U(s_i, \Sigma(s)).$$

Then $f_\xi(s) = \frac{1}{n}\sum_{j=1}^n h_j(s)$. We compute

$$\frac{d}{ds_i}h_i(s) = -U_1(s_i, \Sigma(s)) - U_2(s_i, \Sigma(s)),$$

$$\frac{d}{ds_i}\frac{d}{ds_i}h_i(s) = -U_{11}(s_i, \Sigma(s)) - 2U_{12}(s_i, \Sigma(s)) - U_{22}(s_i, \Sigma(s)),$$

$$\frac{d}{ds_i}\frac{d}{ds_j}h_i(s) = -U_{12}(s_i, \Sigma(s)) - U_{22}(s_i, \Sigma(s)),$$

$$\frac{d}{ds_j}h_i(s) = U_2(\xi, \xi + \Sigma(s_{-i})) - U_2(s_i, \Sigma(s)),$$

$$\frac{d}{ds_j}\frac{d}{ds_k}h_i(s) = U_{22}(\xi, \xi + \Sigma(s_{-i})) - U_{22}(s_i, \Sigma(s)) \text{ for } i \notin \{j, k\}.$$

Then

$$\frac{d}{ds_i}\frac{d}{ds_i}h_i(s) - 2\frac{d}{ds_i}\frac{d}{ds_j}h_i(s) + \frac{d}{ds_j}\frac{d}{ds_j}h_i(s)$$
$$= -U_{11}(s_i, \Sigma(s)) - 2U_{12}(s_i, \Sigma(s)) - U_{22}(s_i, \Sigma(s))$$
$$\quad + 2U_{12}(s_i, \Sigma(s)) + 2U_{22}(s_i, \Sigma(s))$$
$$\quad + U_{22}(\xi, \xi + \Sigma(s_{-i})) - U_{22}(s_i, \Sigma(s))$$
$$= -U_{11}(s_i, \Sigma(s)) + U_{22}(\xi, \xi + \Sigma(s_{-i})),$$

and

$$\frac{d}{ds_i}\frac{d}{ds_i}h_k(s) - 2\frac{d}{ds_i}\frac{d}{ds_j}h_k(s) + \frac{d}{ds_j}\frac{d}{ds_j}h_k(s) = 0 \text{ for } k \notin \{i, j\}.$$

So

$$\frac{d}{ds_i}\frac{d}{ds_i}f_\xi(s) - 2\frac{d}{ds_i}\frac{d}{ds_j}f_\xi(s) + \frac{d}{ds_j}\frac{d}{ds_j}f_\xi(s)$$
$$= \frac{1}{n}\left(-U_{11}(s_i, \Sigma(s)) + U_{22}(\xi, \xi + \Sigma(s_{-i})) - U_{11}(s_j, \Sigma(s)) + U_{22}(\xi, \xi + \Sigma(s_{-j}))\right)$$

---

[9]Note that $\frac{d}{dx}\frac{d}{dx}U(x, nx) = U_{11}(x, nx) + 2nU_{12}(x, nx) + n^2 U_{22}(x, nx)$.

which completes the proof. ∎

We look at some special cases and provide prominent examples for each, thereby uncovering many popular games in which one can beat the play.

Assume first that U is additively separable, so $U(x, y) = g(y) - c(x)$ for some functions c and g. Strategies are strategic complements if g is strictly convex as $U_{12} + U_{22} = g''$. Necessary conditions (Proposition 4) for beating the play are that $c''(\xi) + g''(n\xi) \geq 0$ and $g''(n\xi) \leq \frac{1}{2n}(c''(\xi) + g''(n\xi))$. Sufficient conditions (Proposition 5) are that c and g are convex and $n^2 g''(nx) \leq c''(x)$ for all $x \in A$. So with convex c it is possible to beat the play if goods are strategic complements. A prominent example is a model of private contribution to public goods. Strategies define the contribution, so $A \subseteq \mathbb{R}_+$. $g(y)$ defines the value of a public good with aggregate contribution y. $c(x)$ is the cost of contributing x.

Now assume that $U(x, y) = xg(y) - c(x)$. Necessary conditions for beating the play are $c''(\xi) + xg''(n\xi) \geq 0$ and $g'(n\xi) + \xi g''(n\xi) \leq \frac{1}{2n}(c''(\xi) + \xi g''(n\xi))$. As comparison, following Hahn (1962) and Seade (1979), ξ^n is locally stable (with a continuous time adoption process) if $c''(\xi) > g'(n\xi)$ and $g'(n\xi) + \xi g''(n\xi) < \frac{1}{n}(c''(\xi) - g'(n\xi))$.

The sufficient conditions for beating the play given in Proposition 5 are that c and g are convex and $yg(y)$ is concave on $\mathbb{R}_+$, so $2g'(y) + yg''(y) \leq 0$ for all $y \geq 0$. In particular this means that $g' \leq 0$. So it is possible to beat the play if c and g are convex functions and goods are strategic substitutes (the latter requires $g'(y) + yg''(y) \leq 0$).

Models of contests and fighting (Acemoglu & Jensen, 2013, ch. 5.2) fall in this class. A prominent example is due to Tullock (1980). After transforming the strategic variable, let $g(y) = Vy^{-1}$ with $V > 0$ and $c(x) = x^\alpha$ with $\alpha \geq 1$. Here we set $A = [\varepsilon, \infty)$ for some small $\varepsilon > 0$ to ensure differentiability. Cournot competition with homogeneous goods also falls in this class, as long as demand has no kinks. Strategies are quantities, g captures inverse demand, so $g \geq 0$. c once again captures costs. Sufficient conditions translate into costs being convex, inverse demand being convex and monopoly revenue being concave. For instance, linear inverse demand $g(y) = \max\{0, 1 - y\}$ is a candidate. However, in order to ensure that g is twice differentiable on the relevant domain we need that $A \subseteq [0, \frac{1}{n}]$. The restriction $s_i \leq \frac{1}{n}$

can be justified when c is increasing and $\frac{1}{n} - c\left(\frac{1}{n}\right) \leq -c(0)$ as in this case quantities strictly greater than $\frac{1}{n}$ are strictly dominated. For Cournot competition with kinks in the relevant domain see Section 2.5.

Consider now the slightly more general model in which $U(x,y) = xg\left((1-\gamma)x + \gamma y\right) - c(x,y)$ with $-\frac{1}{n-1} \leq \gamma < 0$ or $0 < \gamma < 1$.

Assume first that $0 < \gamma < 1$. Sufficient conditions are that g is convex, $xg(x)$ is concave and both $c(x,y)$ and $c(x,nx)$ are convex in x. This follows when verifying the following two statements. First of all, $xg\left((1-\gamma)x + \gamma nx\right)$ is concave in x if and only if for any $x \in A$,

$$2\left(1 + \gamma(n-1)\right)g'\left((1-\gamma)x + \gamma nx\right) + x\left(1 + \gamma(n-1)\right)^2 g''\left((1-\gamma)x + \gamma nx\right) \leq 0.$$
(9)

This holds if $xg(x)$ is concave in x as $\gamma > 0$. Second of all, $xg\left((1-\gamma)x + \gamma y\right)$ is concave in x for each y if

$$2(1-\gamma)g'\left((1-\gamma)x + \gamma y\right) + (1-\gamma)^2 xg''\left((1-\gamma)x + \gamma y\right) \leq 0 \text{ for all } x,y \in A \quad (10)$$

which holds if $xg(x)$ is concave in x as $\gamma > 0$. Cournot competition with heterogenous goods falls within this class where g is the inverse demand function and $c(x,y)$ is the cost function that only depends on x. Given the above, sufficient conditions are that costs are convex and demand is convex with monopoly profits being concave.

Now consider the case where $-\frac{1}{n-1} \leq \gamma < 0$. Sufficient conditions are that g is convex, (10) holds and both $c(x,y)$ and $c(x,nx)$ are convex in x. This is because (10) implies (9) when $\gamma < 0$. Note that there are strategic complements if

$$\gamma g'\left((1-\gamma)x + \gamma y\right) + \gamma xg''\left((1-\gamma)x + \gamma y\right) \geq 0,$$

a condition that implies (10) under these restrictions on γ. Bertrand competition with heterogeneous goods and constant marginal costs takes this form (with $g \geq 0$) when considering markups as the strategic variable letting costs $c(x,y)$ be a function of $\left((1-\gamma)x + \gamma y\right)$. Hotelling's (1929) model of spatial competition fits in this specification, setting $\gamma = -1$ and $n = 2$.

Finally, note that the model of price competition with differentiated products of Nocke and Schutz (2018) is an aggregative game in which $U(x,y) = \frac{g(x)}{v+y}$. Sufficient conditions for beating the play are that $g(x)$ and $\frac{g(x)}{v+nx}$ are concave in x.

2.5 Cournot Competition with Kinks

2.5.1 The Analysis

Recall that we considered differentiable payoff functions in the previous section. We now show how to incorporate kinks that arise under Cournot competition when markets do not clear because of excessive quantities. We present the following sufficient conditions.

Proposition 6 *Let $A \subseteq \mathbb{R}_+$ with A convex and compact. Assume $u_i(q) = q_i P(\Sigma(q)) - c(q_i)$ where (i) c is strictly increasing, convex and twice differentiable and (ii) P is nonnegative, decreasing and convex such that P is twice differentiable and $zP(z)$ is concave whenever $P > 0$. Then there exists $\xi \in A$ that beats the play and where ξ is a symmetric NE strategy of Γ^0.*

In particular we find that the Cournot strategy of the game against selves beats the play in the canonical Cournot competition setting with linear inverse demand and constant marginal costs, so where $A = \mathbb{R}_+$, $P(z) = \max\{a - bz, 0\}$ for some $a, b > 0$ and $c(x) = c_0 \cdot x$ for some $c_0 \geq 0$. Note that when $c_0 = 0$ then this game has multiple NE, one that yields a strictly positive payoff for each player and uncountably many others that yield a payoff equal to zero for each player (let each player choose a sufficiently large quantity).

Proof. By the assumptions there exists a pure symmetric NE strategy of Γ^0. Let ξ be such a strategy.

Assume that $P(n\xi) = 0$. Assume that there is some $x \in A$ such that $x < \xi$. As ξ is a symmetric NE strategy of Γ^0, $u(\xi^n) = -c(\xi) \geq u(x, \xi^{n-1}) \geq -c(x)$ which contradicts c being strictly increasing. So $\xi = \inf A$. As P is decreasing we obtain that ξ beats the play.

Assume that $P(n\xi) > 0$. First we show that $f_\xi(q) \geq 0$ when $P(\Sigma(q)) > 0$. As $u(\xi, q_{-i})$ is convex but possibly has a kink when P turns zero we can approximate $u(\xi, q_{-i})$ up to an arbitrarily small error by a smooth convex function (Koliha, 2003). Proposition 5 then shows that f_ξ attains its minimum within the set $\{q : P(\Sigma(q)) > 0\}$ at ξ^n.

Now consider q such that $P(\Sigma(q)) = 0$. As c is convex and $P \geq 0$,

$$f_\xi(q) \geq -c(\xi) + \frac{1}{n}\sum_{i=1}^n c(q_i) \geq -c(\xi) + c\left(\frac{\Sigma(q)}{n}\right).$$

As $P(n\xi) > 0$ and P is decreasing we obtain $\Sigma(q) \geq n\xi$, together with the fact that c is increasing it follows that $f_\xi(q) \geq 0$. ∎

We provide two counterexamples related to Proposition 6, each with $n = 2$. The first example shows that convexity of the inverse demand cannot be dropped from the statement of Proposition 6. Assume that $c = 0$ and $P(z) = \max\{0, 1 - z - \sigma \cdot z^2\}$ for $\sigma > 0$. Then $\xi = \frac{1}{16\sigma}\left(-3 + \sqrt{9 + 32\sigma}\right)$ is a symmetric NE strategy of Γ^0 and easily seen to be the only candidate for beating the play. But ξ does not beat the play as $f_\xi\left(\frac{1}{2\sigma}\left(-1 + \sqrt{1 + 4\sigma}\right) - \xi, 0\right) < 0$ for all $\sigma > 0$.[10]

The second example shows that concavity of $zP(z)$ is cannot be dropped from Proposition 6. Let $P(z) = 26 + c - 28z + 8z^2$ for $z \in \left[0, \frac{3}{2}\right]$ and $P(z) = \frac{1}{z-1} + c$ if $z > \frac{3}{2}$. Then P is twice continuously differentiable, strictly decreasing and convex. However $zP(z)$ is not concave for $z \geq \frac{3}{2}$. The symmetric NE strategy of Γ^0 is given by $\xi = 1$. But ξ does not beat the play as $\lim_{x\to\infty} f_\xi(x,x) = -\frac{1}{2}$.

2.5.2 Performance in Data

In the following we investigate how a strategy that beats the play performs when facing real opponents. The theory tells us that it performs better than when adapting the choice of a random subject. The data will show us how much better it performs and how close it gets to choosing a best response to the empirical distribution of choices. We consider data from laboratory experiments where all subjects have the same monetary payoff function. Although misleading from a conceptual point of view (as beat the play is not about comparing payoffs of different players), the performance is measured by comparing the monetary payoffs of a particular strategy to the average monetary payoffs among the subjects. In particular, we are assuming that player 0 is risk neutral.

---

[10] Nevertheless, the symmetric NE strategy does not perform that bad in this example, its shortcoming is bounded above by 0.0087 for any $\alpha > 0$.

We perform our evaluation using data from a laboratory experiment on Cournot competition with linear demand and no costs where we have found that Cournot output beats the play. Interestingly we did not find any calculations on the performance of the Cournot output in such experiments. In this literature the interest seems to put all focus on how the aggregate demand in the lab differs from the equilibrium aggregate demand. Specifically, we consider the data from the experiments of Huck et al. (2004). Payoffs are given by $u(q) = q_1 \max \{(99 - \Sigma(q)), 0\}$ with action set $A = \{k * 0.01 : k \in \mathbb{N}_0, k * 0.01 \leq 100\}$. The game is played in 25 rounds with fixed partner matching. There were 6 sessions (markets) for each value of $n \in \{2, 3, 4, 5\}$.

For the data analysis we proceed for each value of n (number of players) and round $k \in \{1, ..., 25\}$ as follows. We compute the average payoffs $\bar{\pi}$. Next we replace a random subject in each session and compute the average payoff across all sessions (and all player roles within a session) of this player if replaced player chooses x and denote this by $\pi(x)$. The performance of player 0 as compared to playing like others is then given by $\pi(\xi) - \bar{\pi}$ where ξ is the Cournot output $\frac{99}{n+1}$ (note that $\frac{99}{n+1} \in A$ for all $n \in \{2, ..., 5\}$). This difference is depicted in the figures as a dot. The best possible performance is given by $\max_x \pi(x) - \bar{\pi}$. This difference is depicted in the figures by an inverted triangle. In Figure 1 we plot on the left hand side these two performance measures for the different values of n and k. In the right hand side of this figure we show the average quantities chosen by the subjects, in dependence of n and k.

According to Proposition 6, the dots have to always lie above the value 0, highlighted by a dashed line. By the definitions of beating the play and empirical best response the dots have to lie below the inverted triangle. Interestingly, the dots lie extremely close to the inverted triangle. Choosing the Cournot output leads in this laboratory experiment to payoffs that are extremely close to the maximal possible payoff. We provide a partial explanation. In Proposition 20 in Section D.2 we show that the payoff of the Cournot output is necessarily closer to that of the empirical best response than to the average payoff when $n \leq 5$. In fact we show that the performance of the Cournot output is at least a fraction $\frac{8}{9}$, $\frac{3}{4}$, $\frac{16}{25}$ and $\frac{5}{9}$ of that of the empirical best response for $n = 2$, 3, 4 and 5, respectively.

Note that the observations in this data set come from a repeated game. We do not predict that payoffs of the beating the play strategy maintain its property when
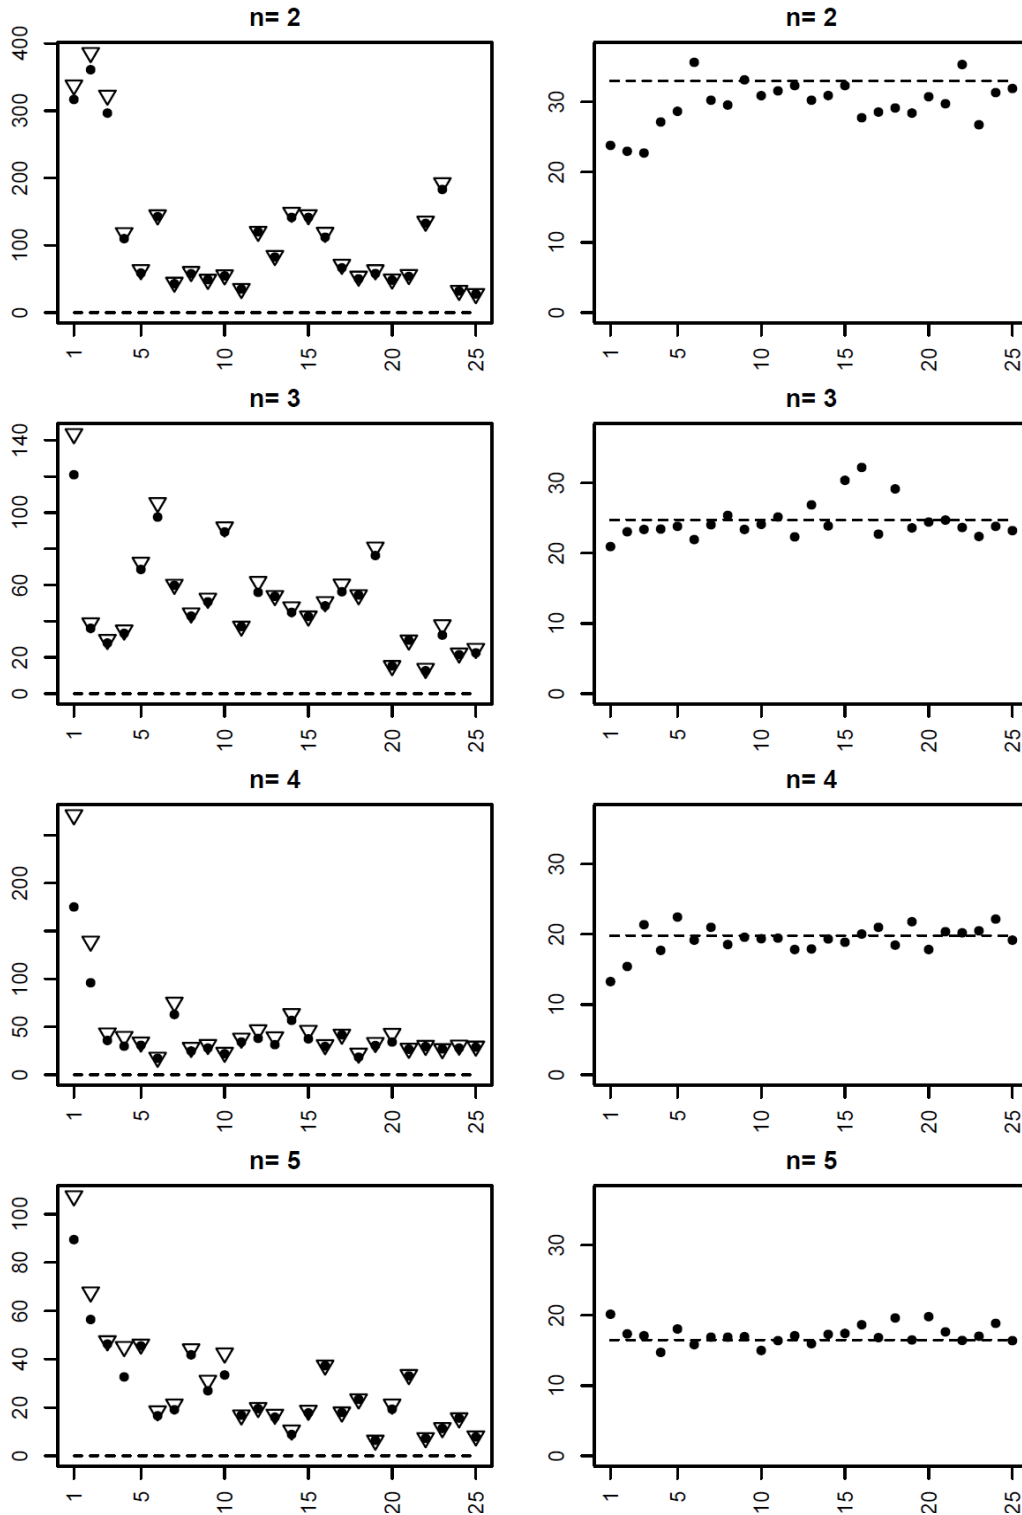
27

Figure 1: Cournot experiment data. Lefthand side panels show the difference to average payoff. Dots show difference from the Cournot output, reversed triangle show difference from the empirical best response. In the righthand side panels, dots show average quantities and dashed line shows the Cournot output level. Panels in different rows show a different number n of players, x achsis in each panel shows the round number.

interacting with the same players over time. For a recommendation for how to play in a repeated game one would have to apply the concept of beating the play to the normal form of the repeated game. We include the data points in rounds $k > 1$ as a proxy for playing a game with others who have had previous experience in the game. Strictly speaking, the performance of the Cournot output in the context of the beating the play concept can only evaluated by looking at the data from round $k = 1$.

3   Extensions

In the following subsections we relax our basic definition in various directions. First we allow a strategy to perform slightly worse than when playing like the others. Next we make the concept more flexible to accommodate that player 0 would never adapt a strictly dominated strategy. Then we include beliefs. We assume that play of others is independent (and identically distributed) and then, more generally, we consider testing performance only on a subset of the strategy profiles.

3.1   Coming Close to Beating the Play

3.1.1   The Concept

Here we select a recommendation for settings where no strategy beats the play. For each strategy we identify the so-called shortcoming. The shortcoming is the maximal amount that the payoff can fall short of that obtained when playing like others. We recommend to use a strategy with the smallest shortcoming. A formal concept is presented that exists under general conditions. Important examples are Bertrand competition with homogeneous goods and first price auctions (the analysis of first price auctions is postponed to Section 3.2).

Definition 4 *(i)* $-1 \cdot \inf_{s \in A^n} f_x(s)$ *is called the* shortcoming *of* $x \in \Delta A$.

(ii) $\xi \in \Delta A$ comes closest to beating the play *if* $\xi \in \arg \max_{x \in \Delta A} \inf_{s \in A^n} f_x(s)$.

(iii) $\xi \in A$ comes closest to beating the play with a pure strategy *if*
$\xi \in \arg \max_{x \in A} \inf_{s \in A^n} f_x(s)$.

It follows that the shortcoming is a non negative number and the shortcoming of a strategy that beats the play is zero. Moreover, a strategy that comes closest to beating the play has the smallest shortcoming.

Remark 2 *We informally say that ξ almost beats the play if its shortcoming is small. While the meaning of small can be context specific, one might categorize small as being an order smaller than $\frac{1}{n} \max_{s \in A^n} u(s)$. An alternative measure for evaluating the size of the shortcoming is provided in Section D.2.*

We proceed by providing general conditions that ensure existence of a strategy that comes closest to beating the play. On the side we identify a useful method for finding such a strategy by characterizing it as part of an equilibrium in a fictitious game against nature.

Let Γ_f be the mixed extension of a two player zero-sum game between yourself and nature in which you choose $x \in A$, nature chooses $s \in A^n$ and your payoff is given by $f_x(s)$.

Proposition 7 *(i) There exists a strategy that comes closest to beating the play if A is finite or if u is continuous and A is a compact subset of $\mathbb{R}^m$.*

(ii) If Γ_f has a NE then ξ comes closest to beating the play if and only if $(\xi, \bar{\sigma})$ is a NE of Γ_f for some $\bar{\sigma} \in \Delta\{A^n\}$.

(iii) If ξ comes closest to beating the play then ξ is not strictly dominated.

Proof. Part (ii) follows directly from the minimax theorem, as a strategy ξ comes closest to beating the play if and only if $f_\xi(\sigma) = \max_{\bar{\xi} \in \Delta A} \min_{\sigma \in \Delta\{A^n\}} f_{\bar{\xi}}(\sigma)$.

Part (i). Following Glicksberg (1952), these assumptions ensure that Γ_f has a NE and so existence follows from part (ii).

Part (iii) is straightforward. ∎

3.1.2   Bertrand Competition Analysis

Consider price competition under homogenous goods where the firm of player 0 has constant marginal costs. We show how to almost beat the play.

Proposition 8 *Let $S = \mathbb{R}_+$ and $c \in \mathbb{R}_+$. Let Q be nonnegative, continuous and decreasing such that $\pi(z) := (z-c)\,Q(z)$ is single peaked with $\{z^*\} = \arg\max_{p\geq 0} \pi(p)$, $z^* > 0$ and $\pi(z^*) > 0$. Assume*

$$u(p) = \begin{cases} \frac{1}{|\{i:p_i=p_1\}|}\pi(p_1) & \text{if } p_1 = \min_i\{p_i\}, \\ 0 & \text{if } p_1 > \min_i\{p_i\}. \end{cases}$$

(i) It is not possible to beat the play.

(ii) The shortcoming of $\xi = c$, the symmetric NE strategy of Γ^0, equals $\frac{1}{n}\pi(z^*)$.

(iii) The unique strategy that comes closest to beating the play with a pure strategy is $\xi \in (c, z^*)$ that solves $(\xi - c)\,Q(\xi) = \frac{1}{n+1}\pi(z^*)$. Its shortcoming is equal to $\frac{1}{n(n+1)}\pi(z^*)$.

The intuition behind the impossibility to beat the play is as follows. Pricing too low means to do worse than playing like others if those choose high prices. Pricing too high means to risk not making any sales and similarly to do worse.

Note that the price that almost beats the play outperforms the symmetric NE strategy of Γ^0 by a factor of $n+1$. It is strictly above marginal costs and is strictly decreasing in n, approaching marginal costs as $n \to \infty$. In particular, the Bertrand paradox does not arise here.

We illustrate with two salient examples. For unit demand and willingness to pay equal to v with $v > c$, the strategy ξ that comes closest to beating the play solves $\xi - c = \frac{1}{n+1}(v-c)$ so $\xi = c + \frac{v-c}{n+1}$. Its shortcoming is equal to $\frac{v-c}{n(n+1)}$.[11] For linear demand $Q(z) = \max\{v - bz, 0\}$ with $b > 0$ and $v > bc$ we solve $(\xi - c)(v - b\xi) = \frac{(v-bc)^2}{4(n+1)b}$ to find $\xi = c + \frac{1}{2b}\left(1 - \sqrt{\frac{n}{n+1}}\right)(v - bc)$. The shortcoming of ξ is $\frac{(v-bc)^2}{4bn(n+1)}$.

Proof. (ii) If $\xi = c$ then

$$\begin{aligned} f_\xi(p) &= -\frac{1}{n}\sum_{i=1}^{n} u(p_i, p_{-i}) = -\frac{1}{n}\left((\min\{p_i\} - c)\cdot Q(\min\{p_i\})\right) \\ &\geq -\frac{1}{n}\pi(z^*) \end{aligned}$$

where equality is obtained if $p_1 = z^*$ and $p_i > p_1$ for all $i > 1$.

---

[11] This can be contrasted to the other two existing models of choice in games without restrictions on behavior of others. Minimax regret prescribes to set a price equal to $c + \frac{v-c}{2}$. Maxmin utility is attained by any price that is greater or equal to c.

(i) As a strategy that beats the play has to be a symmetric NE strategy of Γ^0 and as c is the unique symmetric NE strategy of Γ^0 (Kaplan & Wettstein, 2000), the proof of (i) follows from (ii).

(iii) Clearly, $\xi < z^*$ as the shortcoming of ξ is bounded above by the monopoly profits. Consider only the i-th term $u\left(\xi, p_i\right) - u\left(p_i, p_{-i}\right)$ of $f_\xi\left(p\right)$. There are two potentially worst cases. One is where $p_j = x$ for all j and x is slightly below ξ, in which case $u\left(\xi, p_i\right) - u\left(p_i, p_{-i}\right) \approx -\frac{1}{n}\left(\xi - c\right) Q\left(\xi\right)$. Here we use the fact hat π is single peaked. The other is where $p_j = z^*$ for all j, in which case $u\left(\xi, p_i\right) - u\left(p_i, p_{-i}\right) = \left(\xi - c\right) Q\left(\xi\right) - \frac{1}{n}\left(z^* - c\right) Q\left(z^*\right)$. Both worst cases are independent of i. Hence we choose $\xi < z^*$ to solve

$$\left(\xi - c\right) Q\left(\xi\right) - \frac{1}{n}\pi\left(z^*\right) = -\frac{1}{n}\left(\xi - c\right) Q\left(\xi\right)$$

and hence

$$\left(\xi - c\right) Q\left(\xi\right) = \frac{1}{n+1}\pi\left(z^*\right)$$

which proves the desired claim. Note that a solution to this equation exists as Q is continuous. ∎

Next we show how mixed pricing can improve performance and reduce the short-coming.

Proposition 9 *Consider assumptions on the game against selves specified in Proposition 8. Assume additionally that $\pi'\left(c\right) > 0$ and Q is continuously differentiable.*

(i) There is no single price that comes closest to beating play.

(ii) Let ξ have density $g\left(p\right) = \frac{\pi'\left(p\right)}{n(p-c)}$ for $p \in [z, z^*]$ where $\int_z^{z^*} g\left(p\right) dp = 1$. Then ξ comes closest to beating the play and has a shortcoming equal to $\pi\left(z\right)$.

Proof. Part (ii). We obtain the mixed pricing strategy that comes closest to beating the play by applying Proposition 7 (ii). So we search for a NE of Γ_f, the hypothetical zero sum game between the player 0 choosing ξ and nature who chooses the price vector p of the firms. Let $\pi\left(x\right) = \left(x - c\right) Q\left(x\right)$. Assume that player 0 chooses ξ from a distribution without point masses that has density g. Assume that nature chooses $p = \left(x\right)_{i=1}^n$ for some $x \geq c$. Then

$$\int f_\xi\left(p\right) g\left(\xi\right) d\xi = \int_0^x \left(y - c\right) g\left(y\right) dy - \frac{1}{n}\left(x - c\right) Q\left(x\right).$$

32

Differentiating with respect to x we obtain $(x - c) g(x) = \frac{1}{n} \pi'(x)$. So nature is indifferent if $g(x) = \frac{\pi'(x)}{n(x-c)}$.

Let $z > c$ be such that $\int_z^{z^*} g(p) \, dp = 1$. Note that z exists as $\pi'(x)$ is bounded away from 0 in the neighborhood of $x = c$.

Now we need to find a mixed strategy of nature as a distribution H over x that makes player 0 indifferent on (z, z^*). For $y \in (z, z^*)$ we compute

$$\int f_y(x) \, dH(x) = P_H(x > y) \pi(y) - \frac{1}{n} \int (x - c) Q(x) \, dH(x).$$

To make this expression independent of y we let H have a cdf that satisfies $P_H(x \leq y) = 1 - \frac{\pi(z)}{\pi(y)}$. So H has a point mass $1 - \frac{\pi(z)}{\pi^*}$ at $x = z^*$. It is easily verified that (g, H) is a NE of Γ_f.

Part (i). Assume that the pure strategy y comes closest to beating the play. Then (y, H) must also be a NE of Γ_f. The only candidate for a pure strategy that comes closest to beat the play is presented in Proposition 8, as it is the unique strategy that comes closest to beating the play with a pure strategy. However it is easily seen that this pure strategy is not part of a NE of the zero sum game Γ_f. ∎

We investigate mixed pricing in the two examples presented after Proposition 8. For unit demand we obtain pricing density $g(x) = \frac{1}{n(x-c)}$ on $\left[c + \frac{v-c}{e^n}, v\right]$, and cdf $G(x) = 1 - \frac{1}{n} \ln \frac{v-c}{x-c}$, the associated shortcoming equals $(v - c) e^{-n}$.[12] The use of an appropriate mixed pricing policy yields an improvement as compared to best deterministic price by a factor of $\frac{e^n}{n(n+1)}$, which is equal to 1.23 if $n = 2$ and 4.95 if $n = 5$. For linear demand with $c = 0$ and $b = v = 1$ we obtain pricing density $\frac{1-2x}{nx}$ on $\left[-\frac{1}{2} \text{LambertW}\left(-e^{-(n+1)}\right), \frac{1}{2}\right]$. The associated shortcoming equals $-\frac{1}{4} \text{LambertW}\left(-e^{-n-1}\right)\left(2 + \text{LambertW}\left(-e^{-n-1}\right)\right)$. The magnitude of improvement as compared to the case of pure strategies is slightly larger than it was under unit demand, it is equal to 1.63 for $n = 2$ and 6.72 for $n = 5$.

3.1.3 Bertrand Competition in the Data

We investigate how our strategies that almost beat the play perform in laboratory experiments on Bertrand competition. We are particularly interested in difference

---

[12] This is equivalent to setting the randomized price $p = c + e^{n(z-1)}(v - c)$ where $z \sim U[0, 1]$.

between the single price recommendation (Proposition 8) and the mixed pricing recommendation (Proposition 9).

We consider data from the experiments by Dufwenberg and Gneezy (2000). Specifically, we use the data from treatments 1a-f. There is unit demand, willingness to pay equal to 100, a set of possible prices $A = \{2, .., 100\}$ and number of firms $n = 2, 3, 4$. There are 12 subjects for each treatment. The game was played 10 times with payoffs being generated by randomly matching subjects in each round into groups of size n.

For $n = 2$ the pure strategy recommendation is $\frac{100}{3}$ which is not an element of A. Here we choose the integer that comes closest to beating the play, which is 33. For implementing the mixed strategy recommendation we act as if $A = [0, 100]$.

For the data analysis we randomly match for each n and each round k the subjects into groups of size n. Analogous to our investigation of the Cournot competition data in Section 2.5.2, we plot the performance of the empirical best response (inverse triangle), the pure strategy recommendation ξ_d (dot) and the mixed strategy recommendation ξ_r (circle). We show the results in Figures $2-4$ below. There is a different figure for each value of n, the x axis indicates the round of play. In these figures we include the negative of the shortcoming of each recommendation (small triangle for ξ_d and large for ξ_r) as given by Propositions 8 and 9. By the definition of shortcoming, the circle lies above the large triangle and the dot lies above the small triangle. As pricing strategies are selected to minimize their shortcomings, the shortcoming of the mixed is larger than that of the pure. Hence, the large triangle lies below the small one. Following the definitions of coming close to beating the play and empirical best response, both circle and dot lie below the inverse triangle.

Observe that the strategies that come closest to beating the play perform well above their corresponding theoretical lower bounds (the respective triangles). This is to be expected as these lower bounds are calculated based on very specific worst cases which are unlikely to occur in data. Note also that the presence of these lower bounds determined the specific values for these two strategies. Other strategies reduce these theoretical lower bounds even further. In all cases except for round 2 under $n = 2$, the mixed pricing solution is closer to the average than the pure solution. This is intuitive as the mixed solution tries to equalize losses in many different configurations.

We average the values in each of these figures across the different rounds and

34

n= 2

Figure 2: Bertrand competition with $n = 2$. Difference in payoffs to average payoffs of deterministic solution (dot), randomized solution (circle) and best response to empirical distribution (inversed triangle). Shortcoming of deterministic solution as small triangle and randomized solution as large triangle.

n= 3

Figure 3: Bertrand competition with $n = 3$.

n= 4

Figure 4: Bertrand competition with $n = 4$.

present the results in the following table.

n	2	3	4
pure	2	-1.5	-0.6
mixed	-0.2	-0.9	-0.07
empirical	8.6	3.2	2.2

We observe that the deterministic and mixed pricing rules perform overall very similar in this data set.

Concerning the repeated play in this experiment, the same caveat applies as in our analysis of the Cournot competition data. Strictly speaking, the performance of our recommended strategies can only be evaluated in round 1.

3.2   Adjusting Strategies

In this section we adapt our definition to deal with strategies that player 0 would not choose. For instance, these could be strictly dominated strategies. We do not rule out that she faces such a strategy as we do not know the utilities of others and hence we do not know which strategies are strictly dominated for other players. In Section 3.4 we consider a definition in which some strategy profiles are completely ruled out

36

from the comparison.

To illustrate, consider a bidder in a first price auction with private value v. When making a bid ξ then one might face others that bid more than v but one would never bid more than v. Hence, one would never play like player j within strategy profile s if $s_j > v$. However one might face others making bids above v. We propose how our definitions can be modified.

Assume that player 0 can choose from ΔA but will only choose strategies from a subset $B \subset \Delta A$. Let g be a mapping from ΔA to B with $g(s) = s$ if $s \in B$ where $g(s)$ is the strategy the player chooses if someone recommends s to him. In applications the set B as well as the mapping g are easily identified. For instance, it is natural to have no strictly dominated strategies in B in which case $g(a)$ is an action in B that strictly dominates a. To simplify notation we will drop B from the formal definition, it is as if B is the image of g.

Definition 5 *Let $g : A \to \Delta A$. We say that ξ beats the adjusted play if $g(\xi) = \xi$ and*
$$\frac{1}{n} \sum_{i=1}^{n} \left(u\left(\xi, s_{-i} \right) - u\left(g\left(s_i \right), s_{-i} \right) \right) \geq 0 \text{ holds for all } s \in A^n.$$
We say that ξ comes closest to beating the adjusted play with a pure strategy if $g(\xi) = \xi$ and
$$\xi \in \arg\min_{x \in A} \sup_{s \in A^n} \left\{ \frac{1}{n} \sum_{i=1}^{n} \left(u\left(x, s_{-i} \right) - u\left(g\left(s_i \right), s_{-i} \right) \right) \right\}.$$

We illustrate in a first price auction. Consider player 0 who wishes to bid for a single object in a first price auction with a total of n bidders. Assume that she knows her own value v and regards all player roles as indistinguishable. No assumptions about values or bidding behavior of others are made. An action of a bidder in this game is a bid $b \in \mathbb{R}_+$ where the bidder who makes the highest bid wins the object, ties are broken at random. The payoff u of player 0 when bidding b equals $v - b$ if she wins the object and 0 otherwise. To avoid bidding above v when playing like some bidder who bids above v we assume that the player bids v whenever playing like someone who bid above v. So we set $g(b) = \min\{b, v\}$.

Proposition 10 *Let $g(b) = \min\{b, v\}$. In this first price auction, the bid $\xi = \frac{n}{n+1}v$*

comes closest to beating the adjusted play with a pure strategy. Its shortcoming is

$\frac{v}{n(n+1)}$.

Proof. This result follows from our analysis of Bertrand price competition. Set $p_i = v - b_i$ and $c = 0$ to obtain from Proposition 8 the price $\frac{1}{n+1}v$ and hence the bid ξ is given by $\xi = v - \frac{1}{n+1}v = \frac{n}{n+1}v$. The proof of this result assumed nonnegative prices, which means here that $b_i \leq v$ for all i. It shows that $\xi \in \arg\max_{x \geq 0} \min_{s:s_i \leq v} \left\{ \frac{1}{n} \sum_{i=1}^{n} \left(u\left(x, s_{-i}\right) - u\left(g\left(s_i\right), s_{-i}\right) \right) \right\}$. Note that we can also include profiles s with $s_i > v$ for some i as $\frac{1}{n} \sum_{i=1}^{n} \left(u\left(x, s_{-i}\right) - u\left(g\left(s_i\right), s_{-i}\right) \right) \geq 0$ in such profiles. ∎

3.3 Independent Beliefs

In this section we rule out coordinated and asymmetric play of others.

We illustrate with an example. Consider payoffs such that the game against selves is a Hawk Dove game. Specifically, let $A = \{H, D\}$ where $u\left(D, D\right) = 1$, $u\left(H, D\right) = 2$, $u\left(D, H\right) = 0$ and $u\left(H, H\right) = -1$. Following Proposition 2 it is not possible to beat the play. It is easily shown that the shortcoming of any strategy is equal to $\frac{1}{2}$, attained when the others play $\left(H, D\right)$. Shortcoming is arguably large, so it is not possible to almost beat the play in this game. The inability to almost beat the play rests here on the plausibility of asymmetric profiles $\left(H, D\right)$ and $\left(D, H\right)$.

In this section we assume that player 0 believes that players make their choices independently. As player roles are indistinguishable the circumstances that determine which strategy a player chooses are the same for each player. So player 0 believes she will face identically distributed strategies. In particular, neither $\left(H, D\right)$ nor $\left(D, H\right)$ are candidate environments.

We provide the formal definition.

Definition 6 ξ beats the independent (identical) play *if* $u\left(\xi, \sigma^{n-1}\right) \geq u\left(\sigma^n\right)$ *for all* $\sigma \in \Delta A$.

Note that a strategy that beats the play will also beat the independent play. The following connection to the symmetric Nash equilibria in the game against selves Γ^0 is easy to verify.

Proposition 11 *If ξ beats the independent play then ξ is a symmetric NE strategy of Γ^0.*

We connect to evolutionary game theory. Let Γ^p be the population game in which players choose pure actions from A and payoffs are given by $\bar{u}$ where $\bar{u}(x,y) = u(x, y^{n-1})$ for $x, y \in \Delta A$. In the special case in which $n = 2$ then Γ^p is equivalent to the game against selves Γ^0. We also call Γ^p the corresponding population game.

Remark 3 *A strategy ξ beats the independent play if and only if ξ is a globally neutrally stable strategy (Hofbauer and Sandholm, 2009, HS) of the corresponding population game.*

With this connection we cite HS to obtain the following games in which it is possible to beat the independent play: Hawk Dove game, Rock-Scissor-Paper game in which $w \geq l$ (Example 2.3, HS, see also Footnote 4), a two person war of attrition (Example 2.4, HS) and concave potential games which include congestion games (Sandholm, 2005).

Concerning the observability of the variables used in Definition 6, note that it is easy to estimate σ from data.

3.4 Incorporating Beliefs

Here we consider the setting where player 0 does not conceive each strategy profile as a possible way in which others might play the game. So we only evaluate our criterion for a subset of all strategy profiles. This can be seen as a way to incorporate beliefs, or alternatively, to include a partial understanding of how others play the game. Of course, a strategy that beats the play will maintain its property if some strategy profiles are ruled out. However, other strategies might emerge that also have good properties. Moreover, beating the play might become possible when some strategy profiles are ruled out.

To illustrate, assume that the game against selves is an extremely simple pure coordination game. Specifically, let $A = \{B, C\}$ where $u(a, a') = 1$ if $a = a'$ and $u(a, a') = 0$ if $a \neq a'$. Clearly, it is impossible to beat the play. The unique strategy that comes closest to beating the play is to choose B and C with equal probability,

39

the shortcoming is equal to $\frac{1}{2}$. As the shortcoming is arguably large, it is not possible to almost beat the play.

The inability to almost beat the play in this example is due to the strong externalities. Performance is only good if one can guess what others are doing. The NE concept is based on guessing correctly what others are doing. Here we consider weaker assumptions and assume that player 0 has some idea about how the game is being played. These are not beliefs about what the others are doing but about the possible ways the game is played if one does not take part. Specifically, one identifies a set $U \subset (\Delta A)^n$ of possible strategy combinations and then considers if one can do better than when playing like others for any strategy profiles belonging to this set. A special case of interest is where one believes that others choose similar strategies in which case one can investigate if ξ can beat the play in a neighborhood of ξ^n.

Definition 7 *Let* $U \subset (\Delta A)^n$. *Then* ξ *beats the play in* U *if* $f_\xi(\sigma) \geq 0$ *for all* $\sigma \in U$.

Note that while our motivation is different, beating the play in U generalizes beating the independent play. To see this, simply set $U = \{\sigma^n : \sigma \in \Delta A\}$. Other definitions such as coming closest to beating the play and shortcoming can similarly be adapted to this setting with beliefs.

Returning to the simple pure coordination example above we find that B beats the play if each player is more likely to choose B than C. Note that we can also apply this concept in the Hawk Dove game setting (as described in Section 3.3). There we observe that H beats the play when each player is believed to play D more likely than H.

We extend our insights attained for beating the play as described in Propositions 1 and 2 to this setting. For completeness we also add a straightforward sufficient condition.

Proposition 12 *(i) If* ξ *beats the play in a neighborhood* U *of* ξ^n *then* ξ^n *is a NE of* Γ^0.

(ii) Assume A is finite and that player roles are anonymous. Let $\xi \in \Delta A$. If $u(\xi^k, a^{n-k}) \neq u(\xi^n)$ *for some* $a \in C(\xi)$ *and* $k \in \{1, ..., n-1\}$ *and* ξ *beats the play in a neighborhood U of ξ^n then ξ is a pure strategy.*

(iii) If ξ^n is a strict NE in the game against selves then there is a neighborhood of ξ^n such that ξ beats the play in U.

To illustrate parts (i) and (ii), note that in the Hawk Dove setting (see Section 3.3) there is no strategy ξ that beats the play in a neighborhood of ξ^n. $\xi = \frac{1}{2}[H] + \frac{1}{2}[D]$ is the unique symmetric NE strategy of Γ^0, ξ is not a pure strategy but $u(\xi, D) = \frac{3}{2} \neq \frac{1}{2} = u(\xi, \xi)$.

Proof. Parts (i) and (iii) are easily verified. Part (ii) follows when observing that the proof of Proposition 2 also applies when only considering strategy profiles close to ξ^n. ∎

We connect again to evolutionary game theory by considering beating the independent play in a neighborhood.

Remark 4 *Let A be finite. Combining Definitions 6 and 7, observe that ξ beats the independent play in U for some neighborhood U of ξ^n if and only if ξ is a neutrally stable strategy (Maynard Smith, 1982) of the corresponding population game as defined at the end of Section 3.3. This connects to numerous findings in the evolutionary game theory literature, for example by Banerjee and Weibull (2000) on cheap talk in symmetric 2×2 coordination games.*

Note that the concept of beating the play in a subset of the strategy profiles can be implemented similarly to how it has been suggested by Kasberger and Schlag (2020). Iteratively one reduces the set U until the shortcoming for strategy profiles belonging to U is below some desired level.

4   Conclusion

Equilibrium analysis is pervasive in game theory. In fact, equilibria seem the only natural prediction if players have a mutual understanding that they are trying to solve the game. Yet there are many obstacles to this simple and naive statement. Among the many we note that equilibria need not be unique and that information and objectives of others need not be known, let alone modelable in a satisfactory manor.

We introduce an alternative approach to playing games that does not require the ability to model how others make choices. The concept is easy to compute as demonstrated in this paper with many examples.

One major but simple insight is that only symmetric NE strategies of the symmetric game against selves can beat the (independent) play. Similarly, only Nash equilibria can beat the play under role conditioning.[13] This reveals two novel insights. First of all, to our knowledge, this is the first paper that gives NE a meaning when only the payoffs of a single player are specified. We predict play consistent with an equilibrium when a player has no idea what others will be doing. Second of all, our paper gives symmetric games a new meaning. Traditionally, a symmetric game postulates that all participants are identical, which is hardly ever the case. In our context, players in the symmetric game against selves never play against each other. The symmetric game against selves is only a construct used to interpret a necessary condition. On the side, note that our methodology is not a refinement of the NE concept. A strategy that comes closest to beating the play is often not a symmetric NE strategy of the game against selves.

We uncover a close connection to evolutionary game theory. For instance, the condition for beating the independent play is mathematically equivalent to that of a globally neutrally stable strategy. However, the underlying stories are very different. The one builds on an infinite population of identical subjects comparing payoffs, the other on a single player who is comparing strategies. The different motivations lead to different extensions and variations. Of course the relationship between neutrally stable strategies and symmetric NE strategies is not new. We uncover this finding from a different angle. We benefit from the connection to this evolutionary game theory literature as it enables us to include a few additional examples. Our main characterizations and the bulk of our examples are novel and not present in the literature on neutral stability. One of the main reasons for the emergence of so many results is the much simpler mathematics underlying the concept of beating the play. To beat the play also means to beat the independent play, so results obtained under beating the play can be applied to settings where independent play is more plausible.

---

[13]Similarly, only Bayesian Nash equilibria can beat the play under type conditioning as defined in the online appendix.

We designed our methodology to aid real life choices and to gain insights into incentives in games. It can serve as a benchmark for comparison to peoples' behavior and choices, in the lab or in the field. The key ingredient, a comparison to payoffs of existing strategies, should enhance the perspective of economic agents as opposed to postulating a desiderata. The methodology leads to novel analytic structures that call for new mathematical techniques. This paper offers a glimpse at the possible new insights by presenting many applications and refinements of the basic concept. Exciting extensions for future research include modelling beating the play in extensive form games.

References

[1] Acemoglu, D., and M.K. Jensen (2013), Aggregate Comparative Statics, *Games and Economic Behavior* **81**, 27–49.

[2] Banerjee, A. and J.W. (2000), Weibull Neutrally Stable Outcomes in Cheap-Talk Coordination Games, *Games and Economic Behavior* **32**, 1–24.

[3] Bhaskar, V. (1995), On the Neutral Stability of Mixed Strategies in Asymmetric Contests, *Mathematical Social Sciences* **30**, 273–284.

[4] Camerer, C.F., T.-H. Ho, and J.-K. Chong (2004), A Cognitive Hierarchy Model of Games, *The Quarterly Journal of Economics* **119**, 861–898.

[5] Dufwenberg, M., and U. Gneezy (2000), Price Competition and Market Concentration: An Experimental Study, *International Journal of Industrial Organization* **18**, 7–22.

[6] Glicksberg, I.L. (1952), A Further Generalization of the Kakutani Fixed Point Theorem, with Application to Nash Equilibrium Points, *Proceedings of the American Mathematical Society* **3**, 170–174.

[7] Hahn, F.H. (1962), The Stability of the Cournot Oligopoly Solution, *The Review of Economic Studies* **29**, 329–331

[8] Hofbauer, J. and W.M. Sandholm (2009), Stable Games and Their Dynamics, *Journal of Economic Theory* **144**, 1665–1693.

[9] Hotelling, H. (1929), Stability in Competition, *The Economic Journal* **39**, 41–57.

[10] Huber, P.J. (1981), *Robust Statistics*, John Wiley and Sons, New York.

[11] Huck, S., H.-T. Normann and J. Oechssler (2004), Two are Few and Four are Many: Number Effects in Experimental Oligopolies, *Journal of Economic Behavior and Organization* **53**, 435–446.

[12] Kaplan, T.R, and D. Wettstein (2000), The Possibility of Mixed-strategy Equilibria with Constant-returns-to-scale Technology, *Spanish Economic Review* **2**, 65–71.

[13] under Bertrand competition

[14] Kasberger, B. and K.H. Schlag (2020), *Robust Bidding in First-Price Auctions: How to Bid without Knowing what Others are Doing*, Mimeo, University of Vienna.

[15] Koliha, J.J. (2003), Approximation of Convex Functions, *Real Analysis Exchange* **29**, 465–471.

[16] Linhart, P.B., and R. Radner (1989), Minimax-Regret Strategies for Bargaining over Several Variables, *Journal of Economic Theory* **48**, 152–178.

[17] Maynard Smith, J. (1982), *Evolution and the Theory of Games*, Cambridge University Press, Cambridge.

[18] Nocke, V. and N. Schutz (2018), Multiproduct-Firm Oligopoly: An Aggregative Game Approach, *Econometrica* **86**, 523–557.

[19] Plan, A. (2017), Symmetric n-Player Games, *University of Arizona Economics Working Paper* **17-08**.

[20] Renou, L., and K.H. Schlag (2011), Implementation in Minimax Regret Equilibrium, *Games and Economic Behavior* **71**, 527–533.

[21] Sandholm, W.M. (2005). Negative Externalities and Evolutionary Implementation, *Review of Economic Studies* **72**, 885–915.

[22] Schaffer, M.E. (1988), Evolutionarily Stable Strategies for a Finite Population and a Variable Contest Size, *Journal of Theoretical Biology* **132**, 469–478.

[23] Schlag, K.H. (1998), Why Imitate, and If So, How? A Boundedly Rational Approach to Multi-armed Bandits, *Journal of Economic Theory* **78**, 130–156.

[24] Seade, J. (1979), *The Stability of Cournot Revisited*, Warwick Economic Research Paper **121**.

[25] Selten, R (1980), A Note on Evolutionary Stable Strategies in Asymmetric Animal Conflicts, *Journal of Theoretical Biology* **84**, 93–101.

[26] Stahl, D. O. (1993), Evolution of Smartn Players, *Games and Economic Behavior* **5**, 604–617.

[27] Tullock, G. (1980, Efficient Rent Seeking, in J.M Buchanan, R.D. Tollison and G. Tullock (Eds), *Toward a Theory of the Rent-Seeking Society*, 97–112, College Station, Texas A&M University Press.

[28] von Neumann, J., and O. Morgenstern (1944), *Theory of Games and Economic Behavior*, Princeton, New Jersey: Princeton University Press.

[29] Wilson, R.B. (1987), "Game-Theoretic Analyses of Trading Processes," in Truman Bewley, ed., *Advances in Economic Theory: Fifth World Congress, Econometric Society Monographs*, Cambridge University Press, chapter 2, 33–70.

rest of paper is appendix that is intended for online publication only

A   Extensions that Allow for Asymmetries

We show how one can extend our definitions to the case where all roles in $\{1, ..., n\}$ are no longer indistinguishable. In the first setting we only compare to the strategies of a subset of the players. In the second setting we retain the comparison to all strategies but allow the player to condition her play on the role she takes.

A.1   Subset Comparison

Here we consider a setting where some but not all roles are indistinguishable. Player 0 wants to outperform those that are indistinguishable, regardless of what the other players are doing. To apply this material in this section there have to be at least two roles that are indistinguishable.

We allow the set of pure strategies to depend on the player role. Let A_i be the set of actions that player i chooses from, $i = 1, ..., n$. Indistinguishable player roles necessarily have the same set of actions.

A natural example is a market with buyers and sellers, where a seller tries to beat the play of the sellers, regardless of the behavior of the buyers. To analyze this example in more detail we adapt our definition of coming closest to beating the play. Other definitions are easily adapted too.

Definition 8 *Let $I \subseteq \{1, .., n\}$ with $|I| \geq 2$ such that $A = A_i$ for $i \in I$. Assume that roles in I are indistinguishable. We say that $\xi \in A$ comes closest to beating the play of players belonging to I with a pure strategy if*

$$\xi \in \arg\max_{x \in A} \inf_{s \in \times A_j} \left\{ \frac{1}{|I|} \sum_{i \in I} \left(u\left(x, s_{-i}\right) - u\left(s_i, s_{-i}\right) \right) \right\}.$$

Note that the strategies used by roles that are not indistinguishable influence performance just like some unknown parameter. Player 0 has to anticipate the worst case and tries to do better in this worst case than the strategies used by indistinguishable roles. In Section B we generalize this approach and explicitly allow for unknown parameters.[14]

---

[14]Note that if we let I be a singleton in Definition 8 then we obtain the definition ξ attaining minimax regret.

We use the definition above to make a recommendation to a seller in the following simple market. The good is homogenous. Seller 0 (player 0) produces with marginal costs c_0. There are $n-1$ other potential sellers with $n \geq 2$ and m potential buyers with $m \geq 1$. Each buyer has unit demand and a willingness to pay that is at most $\bar{v}$. The market takes place by each seller choosing a minimal per unit selling price and each buyer a maximal price she is willing to pay for one unit. Trade takes place at the lowest price chosen by the sellers with any buyer who offered a maximal price above this level. Payoffs of seller 0 are given by $p - c_0$ if she sells at price p and 0 otherwise. Then it follows easily that $\xi = c_0 + \frac{\bar{v}-c_0}{n+1}$ comes closest to beating the play of the sellers with a pure strategy. The reason is that the worst cases are attained when all buyers are willing to pay any price below $\bar{v}$. Hence, our results from Proposition 8 apply when setting $Q = m$. In particular, there is no need to explicitly specify a rationing rule.

As a different example consider Cournot competition with a homogenous good. Seller 0 has marginal cost c_0 in a triopoly where one other firm is considered indistinguishable. Let the inverse demand of seller 0 be given by $P(Q) = \max\{1 - Q, 0\}$ and assume $c_0 \in \left[0, \frac{1}{2} - \frac{1}{4}\sqrt{2}\right]$. Then it is easily verified that $\xi = \frac{3}{7} - \frac{1}{7}\sqrt{2} - \frac{1}{7}\left(2\sqrt{2} + 1\right)c_0$ is the best attempt to beat the play of two sellers with a pure strategy. Note that this solution lies below the solution $\frac{1}{4}(1 - c_0)$ obtained when all three sellers are indistinguishable. The associated shortcoming equals $\frac{1}{196}\left(11 - 6\sqrt{2}\right)(1 + 2c)^2$ and ranges from 0.013 to 0.021.

A.2 Heterogeneous Payoffs

We now incorporate heterogeneity within and between player roles in a similar spirit to how it is done in the context of Bayesian games. We maintain the assumption that player roles are indistinguishable. Examples include first price auctions with heterogeneous bidders and Cournot competition with heterogeneous cost functions.

We consider a setting in which player 0 learns her type before the game starts and chooses her strategy prior to learning her type. So the strategy of player 0 is a description of what she chooses for each type. The type determines her set of pure strategies and her payoff function. Types of other players may or may not be

observable.

A.2.1 Private Types

We first consider the case where types of other players are not observable which we refer to as *private type* setting. Player roles are indistinguishable so each player in the game has the same set of possible types and chooses a strategy that conditions on her type.

Note that the type of player 0 cannot be how tall she is as she can hardly choose how to condition on her height before she knows how tall she is. It might be her cost function as a firm if this cost function is linked to recent occurrences such as whole sale prices paid in the morning. It is less plausible model if the cost function is determined by outcomes in the distant past, such as a patent she purchased long ago. First price auctions with private values that are learned before the auction can be a good example.

To help player 0 evaluate her strategy given the play of others prior to the realization of types we introduce a distribution over the joint profile of types. This distribution can be considered the beliefs of player 0. No assumptions on the distribution of the types experienced by other players are made.

We present the formal model. Let T be a set of types with typical representative for the type of player i denoted by θ_i, $i = 1, .., n$. The types of the players in the game is denoted by $\theta = (\theta_1, .., \theta_n)$. Let $\mathcal{T} \subseteq T^n$ be the set of possible vectors of types. At the beginning of the game a vector of types $\theta \in \mathcal{T}$ is drawn from some distribution F. Each player only knows her own type. Let A_{θ_0} be the set of actions of a player with type $\theta_0 \in T$. A pure strategy s_i of player i is mapping from T to $\cup_{\theta_0 \in T} A_{\theta_0}$ such that $s_i(\theta_0) \in A_{\theta_0}$ for $\theta_0 \in T$. Let $\mathcal{S}$ be the set of such pure strategies. Mixed strategies are elements of $\Delta \mathcal{S}$.

Let $u_{(i)}\left(\left(s_j(\theta_j)\right)_j; \theta_i\right)$ be the payoffs of player 0 when she is in the role of player i and player j has type θ_j and chooses pure strategy $s_j(\theta_j)$ for $j \in \{1, ..., n\}$. We extend *weak symmetry* to include types as follows. For each $i \in \{2, ..., n\}$ that there exists a permutation π_i of $\{1, ..., n\}$ with $\pi_i(1) = i$ such that $u_{(i)}(s; \theta_i) = u_{(1)}(s_{\pi_i}; \theta_i)$, $\theta_{\pi_i} \in \mathcal{T}$ and $F(\theta) = F(\theta_{\pi_i})$ for all $\theta \in \mathcal{T}$ where $z_{\pi_i} = \left(z_i, z_{\pi_i(2)}, ..., z_{\pi_i(n)}\right)$. We

simplify notation and write $u\left(s_i, s_{-i}; \theta_i\right)$ instead of $u_{(1)}\left(s_{\pi_i}; \theta_i\right)$. The game against selves is defined by the Bayesian game with pure strategies in $\mathcal{S}$ in which $u_i\left(s; \theta_i\right) = u\left(s_i, s_{-i}; \theta_i\right)$ and type vector θ is drawn according to F.

Note the restrictions on the game that come from indistinguishability. Players with the same type have the same set of pure strategies. Payoffs in the different roles for a given type can be translated into payoffs as perceived as player 1. The joint distribution of types is invariant to the different player role perspectives.

Definition 9 *Assume private types. A strategy $\xi \in \Delta\mathcal{S}$ beats the play under private information if for all $s \in \mathcal{S}$,*

$$\sum_{i=1}^{n} \int \left(u\left(\xi_i\left(\theta_i\right), \left(s_j\left(\theta_j\right)\right)_{j \neq i}; \theta_i \right) - u\left(s_i\left(\theta_i\right), \left(s_j\left(\theta_j\right)\right)_{j \neq i}; \theta_i \right) \right) dF\left(\theta\right) \geq 0.$$

The following implication is immediate.

Proposition 13 *Assume private types. If ξ beats the play under private information then ξ^n is a Bayesian Nash equilibrium of the game against selves.*

We generalize Proposition 5.

Proposition 14 *Assume private types. Let $A \subseteq \mathbb{R}$ be such that A is convex. For each $\theta_0 \in T$ assume that $u\left(s, \theta_0\right) = g_0\left(s_1, \theta_0\right) + s_1 g_1\left(\Sigma\left(s\right), \theta_0\right)$ is such that $g_0\left(\cdot, \theta_0\right)$ is differentiable and concave, $g_1\left(\cdot, \theta_0\right)$ is differentiable and convex such that $z g_1\left(z, \theta_0\right)$ is concave. If $\xi \in A^T$ is a symmetric Bayesian NE strategy of Γ^0 then ξ beats the play under private information.*

Proof. The proof is a straightforward generalization of that of Proposition 5, verifying that f_ξ is convex and that ξ is a local minimum. ∎

Note that the above result does not apply to the typical Cournot competition game in which the inverse demand has a kink. More research is needed to deal with such kinks. It also does not apply first price auctions that have discontinuous payoff functions. An analysis of first price auctions under type conditioning is deferred to future research.

A.2.2 Public Types

We now briefly comment on how the model looks when player 0 can also observe the types of all other players. This is referred to as the *public type* setting. No assumptions are needed on what other players can observe. It is only important that one cannot rule out that all players know the types of all others. We continue to assume that the payoff of player 0 only depends on her own type.

A pure strategy s_i for player i is now a mapping from $\mathcal{T}$ to $\cup_{\theta_0 \in T} A_{\theta_0}$ such that $s_i(\theta) \in A_{\theta_i}$ for all $\theta \in \mathcal{T}$. Let $\mathcal{S}$ be the set of all such pure strategies. In particular, the set of pure strategies for player i also includes those situations in which player i only knows some types of the others. The game that is given by $A_i = A_{\theta_i}$ and $u_i(s) = u(a_i, a_{-i}; \theta_i)$ for $i \in \{1, ..., n\}$ is called the θ *game against selves*. The remaining description and definitions can easily be adapted from the private type setting. Here we talk about *beating the play under public information.*

Consider the special case where there is a single type vector θ. So the support of F consists of n permutations of θ. Then it follows immediately from the definitions that ξ beats the play under public information if and only if

$$\frac{1}{n}\sum_{i=1}^{n} u\left(\xi_i(\theta), a_{-i}; \theta_i\right) \geq \frac{1}{n}\sum_{i=1}^{n} u\left(a_i, a_{-i}; \theta_i\right) \text{ for all } a \in \times A_{\theta_i}. \tag{11}$$

Note that this situation of a single type vector is not mentioned as a plausible configuration, but as a useful building block for compiling sufficient conditions as outlined below.

We obtain some nice connections to NE of asymmetric games when recommending strategies that do not rely on a specific symmetric distribution of types F. The proof is straightforward given the material presented up to now.

Proposition 15 *Assume public types.*

(i) ξ beats the play under public information for each symmetric distribution F over type profiles in $\mathcal{T}$ if and only if (11) holds for all $\theta \in \mathcal{T}$.

(ii) Fix a strategy ξ and a vector of types θ. If (11) holds then $(\xi_i(\theta_i))_{i=1}^{n}$ is a Nash equilibrium of the θ game against selves.

We use part (ii) to make a recommendation for Cournot competition with homogeneous goods and heterogeneous firms. The cost of a firm is her type, let c_i be the

cost of firm i for $i \in \{1, ..., n\}$. Let P be the inverse demand of the market (that may have a kink at $P = 0$). We cannot build on Proposition 6 as its proof relies on the fact that all firms have the same cost function in the game against selves. To simplify the analysis, unlike Proposition 6, we explicitly rule out situations where $q \notin cl\{q : P(s(q)) > 0\}$. Unlike NE analysis we have to otherwise deal with arguable pathological case in which $P = 0$. So we assume that there is some mechanism that ensures that total demand does not reach a zero market price (except on the boundary to the region with positive prices). Specifically, let $q^L = q^L(q)$ be such that $q^L = q$ when $P(s(q)) > 0$ and $q^L \in \partial\{q : P(s(q)) > 0\}$ when $P(s(q)) = 0$. In particular, note that $P\left(s\left(q^L\right)\right) = P(s(q))$.

Proposition 16 *Let $A \subseteq \mathbb{R}_+$ with A convex. Let $\theta = c$ be a vector of cost functions such that c_i is increasing, convex and differentiable for $i \in \{1, ..., n\}$. For $i \in \{1, ..., n\}$ let $u(q_i, q_{-i}; c_i) = q_i^L P\left(s\left(q^L\right)\right) - c_i\left(q_i^L\right)$ where P is nonnegative, decreasing and convex such that if $P > 0$ then P is differentiable and $zP(z)$ is concave. If $\xi(\theta) \gg 0$ is a pure strategy NE of Γ^θ then $\xi(\theta)$ satisfies (11).*

So our recommendation is to identify for each θ a NE of the θ game against selves and to choose the strategy that corresponds to your own type.

Proof. Let $f_\xi(q) := \frac{1}{n}\sum_{i=1}^n u_{(i)}(\xi_i, \sigma_{-i}) - \frac{1}{n}\sum_{i=1}^n u_{(i)}(\sigma_i, \sigma_{-i})$. We compute

$$f_\xi(q) = \frac{1}{n}\sum_{i=1}^n \left(\xi_i P\left(\xi_i + \sum_{j \neq i} q_j\right) - c_i(\xi_i)\right) - \frac{1}{n}\left(\sum_{i=1}^n q_i\left(P\left(\sum_j q_j\right)\right) - c_i(q_i)\right)$$

$$= \frac{1}{n}\sum_{i=1}^n \left(\xi_i P\left(\xi_i + \sum_{j \neq i} q_j\right)\right) - \frac{1}{n}\sum_{i=1}^n c_i(\xi_i) - \frac{1}{n}\left(\sum_{j=1}^n q_j\right)P\left(\sum_{j=1}^n q_j\right) + \frac{1}{n}\sum_{i=1}^n c_i(q_i).$$

Observe that the assumptions made above imply that $f_\xi(q)$ is convex in q when $P(q) > 0$. As P is decreasing, $\{q : P(s(q)) > 0\}$ is a convex set. Hence it is enough to show that $f_\xi(q) \geq f_\xi(\xi)$ when $P(q) > 0$.

Now we show that $f_\xi(q)$ is flat at $q = \xi$ which then implies that ξ is a minimum.

$$\frac{d}{dq_k}f_\xi(q) = \frac{1}{n}\sum_{i \neq k}\left(\xi_i P'\left(\xi_i + \sum_{j \neq i} q_j\right)\right) - \frac{1}{n}\left(\sum_{j=1}^n q_j\right)P'\left(\sum_{j=1}^n q_j\right)$$

$$-\frac{1}{n}P\left(\sum_{j=1}^n q_j\right) + \frac{1}{n}c_k'(q_k)$$

so

$$\frac{d}{dq_k} f_\xi (q) \mid_{q=\xi} = -\frac{1}{n} \xi_k P' \left(\sum_{j=1}^n \xi_j \right) - \frac{1}{n} P \left(\sum_{j=1}^n \xi_j \right) + \frac{1}{n} c'_k (\xi_k) = 0$$

as ξ is a NE of Γ^0. ∎

Let's compare beating the play with and without public information in Cournot competition. Under public information, player 0 evaluates the choice in each role with the corresponding cost function assigned to this role. Performance is then evaluated by taking an average across all roles. Player 0 need not beat the play in each role. To beat the play may mean to take into account low payoffs when having high costs that is offset by high payoffs when having low costs. The solution is to choose the NE of the asymmetric game against selves. In the setting without public information the firm evaluates all circumstances with the same cost function. In some markets she may do worse than the replaced action if this is offset by doing better in the same market in a different role. The solution is to choose a symmetric NE strategy of the game against selves. So the solutions are typically very different and yet both beat the play, in their own context. In the setting without public information, quantities chosen by others are evaluated with own costs. In the public information setting they are evaluated with the cost function associated to the respective role.

B   Uncertainty

In this section we extend our methodology to include uncertainty about parameters of the game. Up to now we only considered uncertainty about the play of others.

Generally speaking, a player often does not know everything about the payoffs in the game she will be playing. We propose how to deal with this uncertainty without introducing priors and do this consistently with our previously introduced methodology. In particular, uncertainty is subjective and only refers to player 0.

Let β be an unknown parameter that describes the details of the game and assume that player 0 knows that β is in a set B. Let $u(s; \beta)$ be the payoff when β is the true parameter. Let

$$f_\xi (s; \beta) := \frac{1}{n} \sum_{i=1}^n (u(\xi, s_{-i}; \beta) - u(s_i, s_{-i}; \beta)).$$

Definition 10 $\xi \in A$ comes closest to beating the play with a pure strategy when β in B if

$$\xi \in \arg\max_{x \in A} \inf_{s \in A^n, \beta \in B} f_x(s; \beta).$$

Note that our concept of coming closest to beating the play of a subset of the players presented in Section A.1 is a special case. Let the parameter β describe the strategies chosen by the players that are not indistinguishable and appropriately define a new game in which n is the number of indistinguishable players in the original game.

We illustrate this concept in Bertrand competition with limited information about the demand. Player 0 is uncertain about the demand and only knows a maximal demand attainable for each price. This scenario might arise when a new firm does not know if it can attract customers, but knows the existing demand. We present a general result and then illustrate the finding with two examples.

Proposition 17 *Consider Bertrand competition as defined in Proposition 8. Let c be the unit cost of production. Let $\bar{Q}(p)$ be the maximal demand at price p, so $Q(p) \leq \bar{Q}(p)$ for all p. Assume that $\bar{Q}$ is continuous and decreasing and $(p-c)\bar{Q}(p)$ is single peaked. Then ξ comes closest to beating the play with a pure strategy when $Q \leq \bar{Q}$ if*

$$\xi \in \arg\max_{x \in [0, p^*]} \inf_{p: p > x} \left\{ -\frac{1}{n}(x - c)\bar{Q}(x), \left((x - c) - \frac{1}{n}(p - c) \right)\bar{Q}(p) \right\}$$

where $\{p^*\} = \arg\max_p \left\{ (p - c)\bar{Q}(p) \right\}$. The associated shortcoming equals $\frac{1}{n}(\xi - c)\bar{Q}(\xi)$.

Proof. Let $\pi^* = (p^* - c)\bar{Q}(p^*)$. We can assume that all firms choose the same price p. Choose some $x \leq p^*$. If $p \leq x$ then

$$f_x \geq -\frac{1}{n}(p - c)Q(p) \geq -\frac{1}{n}(p - c)\bar{Q}(p) \geq -\frac{1}{n}(x - c)\bar{Q}(x).$$

If $p > x$ then $Q(p) \leq Q(x)$ so

$$
\begin{aligned}
f_x &= (x - c)Q(x) - \frac{1}{n}(p - c)Q(p) \geq (x - c)Q(p) - \frac{1}{n}(p - c)Q(p) \\
&\geq \min\left\{ 0, \left(x - c - \frac{1}{n}(p - c) \right)\bar{Q}(p) \right\}.
\end{aligned}
$$

It follows that $\inf_{p:p>\xi} \left\{ \left(\xi - c - \frac{1}{n}(p-c) \right) \bar{Q}(p) \right\} = -\frac{1}{n}(\xi - c)\bar{Q}(\xi) < 0$. Let $\bar{p} \in \arg\inf_{p:p>\xi} \left\{ \left(\xi - c - \frac{1}{n}(p-c) \right) \bar{Q}(p) \right\}$. Observe that the worst case is attained by a demand Q such that Q is constant on $[\xi, \bar{p}]$ which means that the lower bound on f_ξ is tight and the largest possible. Hence, ξ comes closest to beating the play with a pure strategy. ∎

Consider the shortcoming of the strategy that comes closest to beating the play with a pure strategy. By the respective definitions, it is weakly larger when demand can only be bounded from above than when the upper bound is the true demand. In fact, when $\bar{Q}$ is strictly decreasing, then it is strictly larger. This follows when comparing Propositions 17 and 8. This is because the worst case distribution is flat in order to reduce profits at ξ without violating the assumption that Q is decreasing (see proof of Proposition 17). Note also that the recommended strategy is strictly larger under uncertainty. We illustrate when $\bar{Q}(p) = \max\{0, 1-p\}$. Then it follows after some straightforward algebra that

$$\xi = \frac{2c + 2 + cn^2 + (1-c)n - 2(1-c)\sqrt{n}}{4 + n^2}$$

with an associated shortcoming given by

$$\frac{\left(2 + n - 2\sqrt{n}\right)\left(2 + n^2 - n + 2\sqrt{n}\right)}{n\left(4 + n^2\right)^2}(1-c)^2,$$

which for $n \leq 20$ is approximately $\frac{3}{4n(n+4)}(1-c)^2$. If we instead choose the solution for the case where $Q = \bar{Q}$ for the uncertainty setting we find a shortcoming equal to

$$\frac{\left(-n + n\sqrt{\frac{n}{n+1}} + 2\right)^2}{16n}(1-c)^2.$$

We plot shortcomings without the $(1-c)^2$ term, so for the case where $c = 0$, associated to the strategy that comes closest to beating the play with a pure strategy when $Q = \bar{Q}$ (solid line), when $Q \leq \bar{Q}$ (dashed line), and when the solution from $Q = \bar{Q}$ is used for the setting where $Q \leq \bar{Q}$ (dotted line).

On the other hand, if $\bar{Q}$ is derived under unit demand where $\bar{Q}(z) = 1$ for $z \leq v$ and $\bar{Q}(z) = 0$ otherwise, then the solutions of Propositions 8 and 17 coincide.

Figure 5: Shortcoming under Bertrand competition with demand equal to and bounded above by $\bar{Q}(z) = \max\{1 - z, 0\}$.

C Others Try to Beat the Play Too

In this section we show how one can include knowledge that others may be trying to beat the play too. To simplify exposition we assume that all those that are trying to beat the play are identical to player 0. In particular, this means that they have the same utility function as player 0.

Two alternative models come to mind. Under the one there is a probability that other players also attempt to beat the play. In the alternative, it is known that some other players are trying to beat the play too. We illustrate the latter approach. The former approach would be formalized analogous to (Renou & Schlag, 2011).

Assume that player 0 knows that there are k other players who are also trying to beat the play, with $k \in \{1, ..., n - 2\}$. If ξ beats the play under these circumstances, then we assume that all those who are also trying to beat the play also play ξ. This means that we only have to check whether ξ performs better than when playing like others in strategy profiles in which k players are choosing ξ. The formal definition is as follows.

Definition 11 *Let $k \in \{1, ..., n - 2\}$. We say that ξ beats the play with k others*

attempting *if*

$$\frac{1}{n} \sum_{i=1}^{n} \left(u\left(\xi, s_{-i}\right) - u\left(s_i, s_{-i}\right) \right) \geq 0 \ \textit{for all } s \in A^n \textit{ such that } |\{j : s_j = \xi\}| \geq k.$$

Clearly, the definition gets less stringent as k increases. Note that the beating the play condition arises if $k = 0$ is inserted into this definition, the conditions for ξ^n to be a NE of the game against selves emerges if $k = n - 1$. In particular, if a strategy beats the play then it retains its properties if it is known that others are trying to beat the play. Once again it follows that any strategy with this property is a symmetric NE strategy of the game against selves.

We illustrate in Bertrand competition with convex costs and focus on pure strategies. We briefly describe the essentials before making the formal statement. There are k other firms attempting to beat the play. So the prices in the game are given by k firms choosing the candidate price ξ and the other $n-k$ firms choosing some unknown price. When others are also attempting to beat the play then there is no concern of not pricing high enough. There is only pressure to price so low that when the firms with the unknown prices price even lower then player 0's profits when replacing one of them would be negative. For the case of constant marginal costs this means to price at marginal costs and to get zero profits. When costs are strictly convex then pricing sufficiently low works if more than half the players are attempting to beat the play.

Proposition 18 *Consider Bertrand competition with a homogeneous divisible good with $n \geq 3$, $k \in \{1, .., n-2\}$ $S = \mathbb{R}_+$, increasing and convex costs with $c(0) = 0$, continuous and decreasing demand Q and payoffs given by*

$$u(p) = \begin{cases} p_1 \frac{1}{|\{i:p_i = p_1\}|} Q(p_1) - c\left(\frac{1}{|\{i:p_i = p_1\}|} Q(p_1) \right) \ \textit{if } p_1 = \min_i \{p_i\}, \\ 0 \ \textit{if } p_1 > \min_i \{p_i\}, \end{cases}$$

with $u\left(x^i, \infty^{n-i}\right)$ as function of x being single peaked for each $i \in \{1, .., n\}$.[15]

(i) It is not possible to beat the play with a pure strategy.

(ii) If (a) marginal costs $c'(Q)$ are constant, (b) $n \leq 4$ or (c) $k \geq \frac{n}{2} - 1$ then it is possible to beat the play with k others attempting. Under (a) choose $\xi = c'$. Under

---

[15]$u\left(x^i, \infty^{n-i}\right)$ stands for $u\left(x^i, y^{n-i}\right)$ for some $y > x$.

(b) and (c) choose $\xi \geq 0$ such that $u\left(\xi^{k+1}, \infty^{n-k-1}\right) = 0$ or any $\xi \geq 0$ that solves

$$u\left(\xi^{k+1}, \infty^{n-k-1}\right) \geq 0 \geq u\left(\xi^{n-k-1}, \infty^{k+1}\right).$$

(iii) If $n \geq 5$, $k < \frac{n}{2} - 1$ and costs are strictly convex then it is not possible to beat the play with a pure strategy with k others attempting.

Note that ξ is a symmetric NE strategy if and only if $u\left(\xi^n\right) \geq \max\left\{0, u\left(\xi, \infty^{n-1}\right)\right\}.$

Proof. Part (i). Assume that $\xi \in A$ beats the play. Let $z > 0$ be such that $u\left(z^n\right) > 0$. Following the arguments made in the proof of Proposition 8, replacing z^* by z, we obtain

$$\inf_{p \in \mathbb{R}_+^n} f_\xi\left(p\right) \leq \min\left\{-u\left(\xi^n\right), u\left(\xi, \infty^{n-1}\right) - u\left(z^n\right)\right\}.$$

So if ξ beats the play then $u\left(\xi^n\right) = 0$ and $u\left(\xi, \infty^{n-1}\right) \geq u\left(z^n\right)$ but $u\left(\xi, \infty^{n-1}\right) \leq nu\left(\xi^n\right) = 0$ which is a contradiction to $u\left(z^n\right) > 0$.

Parts (ii) and (iii). Assume that $\xi \in A$ beats the play with k others are attempting. We are interested in $\inf_p f_\xi\left(\xi^k, p_{k+1}, ..., p_n\right)$. There are only two potential worst cases. Consider the situation in which the $n - k$ other firms price higher than ξ. Then ξ needs to satisfy

$$u\left(\xi^{k+1}, \infty^{n-k-1}\right) = \xi \frac{1}{k+1} Q\left(\xi\right) - c\left(\frac{1}{k+1} Q\left(\xi\right)\right) \geq 0 \tag{12}$$

as the $n - k$ other firms are making no sales. Consider the situation in which the $n - k$ other firms choose the same price which is slightly below ξ. Then ξ needs to satisfy

$$0 \geq u\left(\xi^{n-k-1}, \infty^{k+1}\right) = \xi \frac{1}{n-k-1} Q\left(\xi\right) - c\left(\frac{1}{n-k-1} Q\left(\xi\right)\right). \tag{13}$$

Here we use continuity of Q and single peakedness. It follows that $\xi \in A$ beats the play with k others attempting if and only if (12) and (13) hold.

We show that there exists ξ such that $u\left(\xi^{k+1}, \infty^{n-k-1}\right) = 0$. As $u\left(x^{k+1}, \infty^{n-k-1}\right)$ is continuous in x and $u\left(0^{k+1}, \infty^{n-k-1}\right) \leq 0$ we only have to find y such that $u\left(y^{k+1}, \infty^{n-k-1}\right) > 0$. Note that $u\left(x^{k+1}, \infty^{n-k-1}\right) = x\frac{1}{k+1}Q\left(x\right) - c\left(\frac{1}{k+1}Q\left(x\right)\right) \geq \frac{1}{k+1}Q\left(x\right)\left(x - \frac{c(Q(x))}{Q(x)}\right)$ where $\frac{c(Q(x))}{Q(x)}$ is decreasing when $Q\left(x\right) > 0$. As $u\left(x^{k+1}, \infty^{n-k-1}\right)$ is single peaked we have $Q\left(0\right) > 0$. Then if $y > \frac{c(Q(0))}{Q(0)}$ then $u\left(y^{k+1}, \infty^{n-k-1}\right) \geq \frac{1}{k+1}Q\left(y\right)\left(y - \frac{c(Q(y))}{Q(y)}\right) \geq \frac{1}{k+1}Q\left(y\right)\left(\xi - \frac{c(Q(0))}{Q(0)}\right) > 0$.

Assume that $k + 1 \geq n - k - 1$. We show that (12) and (13) are compatible with each other. In particular, this means that if ξ solves $u\left(\xi^{k+1}, \infty^{n-k-1}\right) = 0$ then (12) and (13) hold.

By convexity of c we have $\frac{k+1}{n-k-1} c\left(\frac{1}{k+1} Q\left(\xi\right)\right) \leq c\left(\frac{1}{n-k-1} Q\left(\xi\right)\right)$, so

$$(n - k) u\left(\xi^{n-k-1}, \infty^{k+1}\right) \leq (k + 1) u\left(\xi^{k+1}, \infty^{n-k-1}\right)$$

and hence

$$u\left(\xi^{k+1}, \infty^{n-k-1}\right) \geq \frac{n - k - 1}{k + 1} u\left(\xi^{n-k-1}, \infty^{k+1}\right).$$

Now consider $k + 1 < n - k$. Then the arguments above show that

$$u\left(\xi^{k+1}, \infty^{n-k-1}\right) \leq \frac{n - k - 1}{k + 1} u\left(\xi^{n-k-1}, \infty^{k+1}\right)$$

with strict inequality holding whenever c is strictly convex. So if c is strictly convex then we find that (12) and (13) cannot both be true. ∎

D   Alternative Definitions

D.1   Beating the Loser

Consider a weaker definition than beating the play in which the player wishes to choose a strategy that does better than some strategy used by others. So instead of doing better than a random strategy the objective is now to do better than the worst strategy. To have this property seems to be a minimal requirement. Clearly, any strategy that does well when comparing to strategies used by others should also outperform the worst strategy used by others. However, it is doubtful that this property alone can be used to justify a particular strategy. We acknowledge that the definition can be useful to understand what is going on when one cannot beat the play. However the insights are not substantial and hence we have placed this material in the appendix.

Definition 12 *The strategy ξ beats the loser if for each $s \in A^n$ there exists $i \in \{1, .., n\}$ such that $u\left(\xi, s_{-i}\right) \geq u\left(s_i, s_{-i}\right).$*

Note that the definition remains unchanged if instead one formulates it with respect to mixed strategies, so requires for each $\sigma \in \times (\Delta A)^n$ that there exists i such that $u(\xi, \sigma_{-i}) \geq u(\sigma_i, \sigma_{-i})$.

Proposition 19 *(i) If ξ beats the loser then ξ is a symmetric NE strategy of Γ^0.*

(ii) If ξ beats the play then ξ beats the loser.

(iii) If ξ does not beat the play and $\inf_{s \in A^n} f_\xi(s) = \inf_{a \in A} f_\xi(a^n)$ then ξ does not beat the loser.

Proof. Parts (ii) and (iii) follow from the definitions.

For part (i), assume that $u(z, \xi^{n-1}) > u(\xi^n)$ for some $z \in A$. Let $\sigma = (1 - \varepsilon)\xi + \varepsilon z$. Then $u(\sigma, \sigma^{n-1}) > u(\xi, \sigma^{n-1})$ if ε is sufficiently small which implies that ξ does not beat the loser. ■

For instance, from the above and our analysis in Section 3.1.2 it follows that it is not possible to beat the loser in Bertrand competition as described in Proposition 8.

The following example shows that not every strategy that beats the loser also beats the play.

Example 1 *Let u be payoffs as in a Rock Scissors Paper game, so $A = \{R, S, P\}$ such that $u(R, S) = u(S, P) = u(P, R) = w > 0$ and $u(S, R) = u(P, S) = u(R, P) = -l < 0$. The symmetric NE strategy ξ of Γ^0 puts equal weight on each of the pure strategies. It is easily verified that ξ beats the loser if and only if $w \geq l$ while ξ beats the play if and only if $w = l$.*

D.2 The Magnitude of Beating the Play

For a strategy that beats the play it can be interesting to capture where its performance lies between playing a best response and realizing payoffs as if playing like others. We do this as follows.

Definition 13 *ξ beats the play by magnitude λ^* if λ^* is the largest value of $\lambda \leq 1$ that solves*

$$\frac{1}{n} \sum_{i=1}^{n} u(\xi, s_{-i}) \geq \lambda \frac{1}{n} \max_{x \in A} \left\{ \sum_{i=1}^{n} u(x, s_{-i}) \right\} + (1 - \lambda) \frac{1}{n} \sum_{i=1}^{n} u(s_i, s_{-i}) \text{ for all } s \in A^n.$$

We present a formula for λ^*. Let $M = \left\{ s \in A^n : \max_{x \in A} \left\{ \sum_{i=1}^n u(x, s_{-i}) \right\} > \sum_{i=1}^n u(s_i, s_{-i}) \right\}$. If $M = \emptyset$ then $\lambda^* = 1$. If $M \neq \emptyset$ then

$$\lambda^* = \min_{s \in M} \frac{\sum_{i=1}^n u(\xi, s_{-i}) - \sum_{i=1}^n u(s_i, s_{-i})}{\max_{x \in A} \left\{ \sum_{i=1}^n u(x, s_{-i}) \right\} - \sum_{i=1}^n u(s_i, s_{-i})}. \tag{14}$$

We illustrate in Cournot competition with linear demand and constant marginal costs. We impose that quantities are never excessive and use the notation from Section A.2.2.

Proposition 20 *Let $A = [0, 1]$, $0 \leq c_0 < 1$ and $u(q) = q_1^L \left(1 - \Sigma \left(q^L \right) \right) - c_0 \cdot q_1$. Then $\xi = \frac{1 - c_0}{n+1}$ beats the play by magnitude*

$$\lambda^* = \frac{4n}{(n+1)^2}.$$

The values for λ^* for $n = 2, 3, 4, 5$ are given by $\frac{8}{9}, \frac{3}{4}, \frac{16}{25}$ and $\frac{5}{9}$ respectively. Note that the proof reveals that the right hand side in (14) is independent of q if $\Sigma(q) \leq 1$ and $\xi + \Sigma(q_{-i}) \leq 1$ for all i.

Proof. Consider q such that $\Sigma(q) \leq 1$. We verify

$$\frac{1}{n} \sum_{i=1}^n u(q_i, q_{-i}) = \frac{1}{n} \Sigma(q)(1 - \Sigma(q)),$$

$$\frac{1}{n} \sum_{i=1}^n u(\xi, s_{-i}) \geq \xi \left(1 - \xi - \Sigma(q) \left(1 - \frac{1}{n} \right) \right),$$

$$\max_{x \in A} \sum_{i=1}^n u(x, q_{-i}) = z \left(1 - z - \Sigma(q) \left(1 - \frac{1}{n} \right) \right) \text{ with } z = \frac{1}{n+1} - \frac{n-1}{2} \left(\frac{\Sigma(q)}{n} - \xi \right).$$

The rest of the proof is straightforward. ∎

One might also wish to compute magnitudes for strategies that cannot beat the play, in this case $\lambda^* < 0$. For example, consider Bertrand competition under unit demand, willingness to pay v and no costs. Then the strategy that comes closest to beating the play is $\frac{v}{n+1}$, its shortcoming is $\frac{v}{n(n+1)}$. For symmetric allocations and own price z we obtain

$$\inf_{s \in A^n} \frac{\sum_{i=1}^n u(z, s_{-i}) - \sum_{i=1}^n u(s_i, s_{-i})}{\max_{x \in A} \sum_{i=1}^n u(x, s_{-i}) - \sum_{i=1}^n u(s_i, s_{-i})} = \inf_{y < a \leq v} \left\{ -\frac{1}{n-1}, \frac{nz - v}{(n-1)v} \right\} = -\frac{1}{n-1}.$$

Consider now independent play as modeled in Section 3.3. We analogously define the magnitude λ^* of beating the independent play by the largest $\lambda \in [0, 1]$ that solves

$$u\left(\xi, \sigma^{n-1} \right) \geq \lambda \max_{x \in A} u\left(x, \sigma^{n-1} \right) + (1 - \lambda) u(\sigma^n) \text{ for all } \sigma \in \Delta A.$$

So if $M_0 = \{\sigma \in \Delta A : \max_{x \in A} u(x, \sigma^{n-1}) > u(\sigma^n)\} \neq \emptyset$ then

$$\lambda^* = \min_{\sigma \in M_0} \frac{u(\xi, \sigma^{n-1}) - u(\sigma^n)}{\max_{x \in A} u(x, \sigma^{n-1}) - u(\sigma^n)}.$$

We verify that the magnitude of beating the independent play is equal to 0 in the Hawk Dove game setting described in Section 3.3.