Partial Exposure in Large Games*

Working Draft, Comments are Welcome

Ronen Gradwohl † Omer Reingold ‡

September 3, 2009

Abstract

In this work we introduce the notion of partial exposure, in which the players of a simultaneous-move Bayesian game are exposed to the realized types and chosen actions of a subset of the other players. We show that in *any* large simultaneous-move game, each player has very little regret even after being partially exposed to other players. If players are given the opportunity to be exposed to others at the expense of a small decrease in utility, players will decline this opportunity, and the original Nash equilibria of the game will survive.

Keywords: Noncooperative games, large games, Nash equilibrium, ex post Nash.
JEL classification: C72.

1 Introduction

In a recent series of works, Kalai (2004, 2005) initiated a study of the properties of *general* strategic noncooperative games with a large number of players. His results uncovered the following properties: The equilibria of large simultaneous-move Bayesian games are ex post Nash and structurally robust. The ex post Nash property means that even if players see the realized types and chosen actions of other players, they have little regret. In other words, their chosen strategy is still an (almost) best response, and they can not benefit much from any unilateral change of action. Structural robustness means that the strategies form an equilibrium even in extensive versions of the game, which can include alterations such as sequential play, possibility of revision, the addition of players, and more. The results are very meaningful in the sense that they eliminate many modelling difficulties associated with the structure of the game.

However, these results do depend on two conditions that limit their generality:

*We would like to thank Ehud Kalai and Yaron Azrieli for interesting and useful conversations. We are grateful to Ariel Yadin and Amir Yehudayoff for helpful discussions throughout the course of this research. We also thank seminar participants at Northwestern University and at the Technion. Finally, we are grateful to the anonymous referees for valuable input.

†Department of Computer Science and Applied Mathematics, The Weizmann Institute of Science, Rehovot, 76100 Israel. E-mail: ronen.gradwohl@weizmann.ac.il.

‡Department of Computer Science and Applied Mathematics, The Weizmann Institute of Science, Rehovot, 76100 Israel. E-mail: omer.reingold@weizmann.ac.il. Research supported by US-Israel Binational Science Foundation Grants 2002246 and 2006060.

1. The game is semi-anonymous and uniformly equicontinuous. Roughly, semi-anonymity means that players' utility functions do not depend on the specific types and actions of players, but only on the fraction of players that obtained each type-action pair. Equicontinuity means that small changes in these fractional values do not affect players' payoffs too much.

2. The types of the players are mutually independent.

In this work we will examine the first condition, argue that it is desirable to weaken it, and attempt to do so. In the companion paper to this work (Gradwohl and Reingold 2007) we address the second condition.

We now turn our attention to the anonymity and continuity conditions on the game. Kalai (2004) showed that these are sufficient conditions for ex post Nash. But are they necessary? It turns out that there are many non-continuous games in which the equilibria are not ex post Nash or structurally robust (see Section 1.1 for an example). So on the one hand, interactions of many agents are potentially highly-sensitive to alterations in the order of play and the information structure. On the other hand, often such specifics of the game are not available to either the modeler or even the players themselves. For example, a game that models the interaction of many agents on the Internet or some other distributed setting can not be sensitive to these issues, since synchrony and the absence of faults can not be guaranteed here (Halpern 2003). Thus, perhaps many of the phenomena we wish to model have some sort of innate robustness.

How can we capture this apparent robustness game-theoretically, without being restricted to the limitations of continuous games? We do this by weakening the ex post Nash requirement, and then showing that this weaker property holds for a much more general family of games. Perhaps surprisingly, the only restriction we place on large games is that the types be independent. In the companion paper to this work, we weaken even this requirement.

We weaken the ex post Nash property by only allowing *partial exposure*. Instead of being exposed to the realized types and chosen actions of all other players, each player only observes the values of some of the players. Why is this a reasonable assumption? Recall that we are interested in large games. Perhaps we are modelling some interaction of agents on the Internet, or perhaps we are analyzing some large market game. In such a situation, it may simply not be feasible for a player to gather information about all other players. Or perhaps gathering information has a cost, and each player is only willing to expend a certain amount of resources for this task.

We will argue that, as a modelling tool, the notion of partial exposure is both a reasonable and a useful assumption. Additionally, we view our result as a meaningful characterization of a property of large games. To the best of our knowledge, this is the only known property of general large noncooperative games that does not necessarily assume anonymity or continuity. Prior work on large noncooperative games, aside from the work of Kalai (2004, 2005), has included issues such as purification – see Shmeidler (1973) and subsequent works, such as Cartwright and Wooders (2009), Carmona (2008), and Carmona and Podczeck (2009) – and the properties of equilibrium distributions (Blonski (2000)). However, all the results on the former assume continuity, and results on the latter only apply to games that are anonymous. Additional related work in this literature includes

Azrieli (2009) and Wooders, Cartwright, and Selten (2006), and these also assume a form of continuity.

Our results, on the other hand, are extremely general. We show that for **every** large simultaneous-move game, each player has little regret with high probability after being exposed to some of the other players' realized types and chosen actions.

Stability under costly exposure: An interesting interpretation of our result is the following. Suppose the players play the game, and each is then given an opportunity to be exposed to the realized types and chosen actions of some of the other players, after which she may revise her own strategy. However, this exposure comes at a *small* cost (in terms of utility). Then our theorem states that players will **not want** to be exposed, because the cost will outweigh the possible benefit. Hence, the original equilibria of the game will survive.

As an example, such a situation can arise in a game in which players are permitted to obtain a survey of others' types and actions, in exchange for a small fee, and then modify their own action. In this case, the survey is the exposure, and the fee is the small decrease in utility. Our result implies that a best-response for players is to decline the opportunity to conduct this survey.

The remaining issue is what it means to see the types and actions of **some** other players. One possible interpretation is that players are exposed to values of a random subset of the other players. Alternatively, each player is exposed to each other player independently with some probability. We show that in this scenario of random exposure, players will not have much regret.

Can we strengthen this to allow exposure to an arbitrary subset of the other players? Later we will show that in general large games this is unattainable. Intuitively, if a player's utility depends only on one other player, then exposure to that other player can be very beneficial. However, we still prove a general theorem that, in a sense, shows that this example is a worst case. For each player, we "forbid" a (small) subset of the players, but then allow the player to choose an arbitrary subset of the remaining players to which she will be exposed. Intuitively, the forbidden players contain some of the more influential ones. Suppose, for example, that the game is played between financial entities. A player may be an individual investor, who has little global influence on the game, or a large corporation with larger influence. An individual may travel around, meet other individuals, and be exposed to their types and actions. But she will not be able to obtain knowledge of the types and actions of the large corporations.

The setting of exposure to a random subset, on the other hand, corresponds more closely to a game played on the Internet, in which players may obtain information about other players at random.

We note also that our results hold for all games, but are meaningful only for large games. For example, in the random exposure setting we will show that there is some positive number β such that for all n the following holds: a player in an n-person game will not want to change her strategy even if she is randomly exposed to the realized types and chosen actions of βn other players. Of course, if $n < 1/\beta$, then this statement is trivially true, since the player is not exposed to any other player.

Finally, we note that this paper is mostly motivated by two works – Kalai (2004) and Al-Najjar and Smorodinsky (2000). Our primary motivation, as well as the model with which we work, are inspired by Kalai (2004). Our techniques and the proofs of our theorems rely on notions of the influence of variables in functions, most notably the results of Al-Najjar and Smorodinsky (2000).

In summary, this work examines several flavors of partial exposure, which correspond to natural phenomena we wish to model. Hence, partial exposure can be viewed as a game-theoretic explanation and modelling tool for robustness properties of large games.

Organization The rest of the paper is organized as follows. The remainder of this section contains examples (Section 1.1) and some further discussion (Section 1.2), and in Section 2 we describe the model and some necessary definitions. In Section 3 we state and discuss our main theorems. Section 4 contains the proofs of our theorems, and Section 5 provides examples of games that demonstrate the tightness of our results.

1.1 Examples

In order to provide further motivation and intuition, we examine an example – the "Computer Choice Game" of Kalai (2004).

Example 1.1 There are n players, each of which has to choose between a computer I and a computer M. In addition, each player is equally likely to be a type who likes I and a type who likes M. The payoffs for each player come both from agreement with other players, and from agreement with one's own preference. Specifically, the payoff for each player is 0.9 times the proportion of other players her choice matches, plus 0.1 if she matches her own type. If each player knows her own realized type, then there are three Nash equilibria in this game: all players choose I, all players choose M, and each player chooses her preferred type.

The constant strategies are clearly ex post Nash, since, even after seeing the choices of other players, no player has any incentive to change her own choice. Actually, no player even needs to see the choices of other players, since she knows what they will be. Furthermore, these strategies are robust in the following sense: even if the players do not play the game simultaneously, but rather sequentially and even with possibility of revision, the original strategies are still best-response. However, the same is not immediately true for the choose-your-favorite-computer strategies. Sometimes it may be preferable for some player to match the choice of another who has already played. Nevertheless, Kalai (2004) showed that this is not likely to occur: with very high probability, no player has much regret. In addition, the original strategy is still an approximate-Nash equilibrium in any extensive version of the game (playing sequentially and with possibility of revision).

We now modify the game slightly.

Example 1.2 Suppose the scenario is the same as in the Computer-Choice Game, but with slightly different payoffs. Each player still gets 0.1 if she matches her own preferred type. However, she gets an additional 0.9 if she matches the **majority** of other players. Again, there are three Nash equilibria in this game: the two constant strategies and the choose-your-favorite-computer strategy.

In this example, the choose-your-favorite-computer strategy is no longer ex post Nash. In fact, after seeing the choices of other players, about half of the players will want to change their own strategy. So what went wrong? The answer is that the game no longer satisfies the uniform equicontinuity condition.

Intuitively, however, it seems like this strategy should be robust somehow. Suppose we want to buy a computer. We interact with many other potential computer buyers, and learn about their preferences and their eventual choices. Ultimately, however, our choice is probably close to what it was originally.

The point to notice is that if we want to buy a computer, we can not know what all other players like and what they buy. Perhaps a player goes to a computer store, where she sees and interacts with some other players, and learns about their preferences and their eventual choices. But certainly not all players go to the store at the same time, nor do they interact with or learn about all other players.

Our results show that if a player is exposed to the types and actions of only a fraction of the players, then she does not have too much regret with high probability. In other words, if players play the choose-your-favorite-computer strategy, then even after being exposed to the preferences and choices of some other players, they have little incentive to change their own choice. In fact, for this game each player can be exposed to an arbitrary subset of other players of size linear in the number of players.

Now consider an additional variant of the Computer-Choice game, called the Modern-Day-Computer-Choice game.

Example 1.3 Suppose the scenario is the same as in Example 1.2, except that players purchase their computers online, over some insecure network. Because of the lack of security in the network, some information may leak: for any pair of buyers, there is some small probability q, independent of all other events, that one of the buyers observes the other's purchasing order. Players do not know q, but rather some upper bound on its value.

In this example, players may wish to delay their purchase as long as possible, and attempt to obtain as much information about other players as they can. The question is, will the information they eventually obtain help them get a larger payoff? Our results imply that, for sufficiently small $q > 0$, players can just as well purchase their computers immediately. The potential information they will obtain will most likely not cause them to change their strategy.

1.2 Perspective and Further Directions

We believe that a main contribution of this paper is the introduction of partial exposure to capture a form of robustness that is weaker than the structural robustness of Kalai (2004). In this paper our aim was to obtain a result that is as general as possible, and in fact the examples of Section 5 show that our results are tight. One direction for future work is to examine a more restricted setting and provide stronger robustness results. The result of Kalai (2004), for example, is much stronger than ours, albeit for the restricted class of games that are semi-anonymous and uniformly equicontinuous. A subsequent work of Gradwohl and Reingold (2008) examines games that are semi-anonymous but *not* equicontinuous, and also obtains stronger results.

A different extension due to Gradwohl and Reingold (2008) is to place partial exposure alongside other types of "faults" that can occur when games are played over a network, and examine their joint robustness properties.

2 The Model and Some Definitions

In general we will denote sets by capital letters, random variables that take the values of elements in the sets by bolded capital letters, and elements of sets by the corresponding lower case letters. For example, X is some set, $x \in X$ is an element of X, and $\mathbf{X}$ is a random variable that takes values in X.

We use the standard definition of a Bayesian game (see Kalai (2004)). Let $\mathcal{T}$ and $\mathcal{A}$ be sets representing possible player types and actions. Unless noted otherwise, for all games discussed in this paper, $\mathcal{A}$ is finite but $\mathcal{T}$ may be unbounded (but countable). Some of our results are meaningful only when the type space is finite, and when this is the case it will be explicitly noted. Denote by $\mathcal{C} \equiv \mathcal{T} \times \mathcal{A}$ as the set of possible type-action characters of players. We now define the family of games $\Gamma(\mathcal{C})$ as follows.

Definition 2.1 (game) *A game $G \in \Gamma(\mathcal{C})$ is described by a five-tuple $G = (N, T, \mathbf{T}, A, u)$ as follows:*

- *$N = \{1, \ldots, n\}$ is the set of players.*
- *$T = T_1 \times \ldots \times T_n$ is the set of player types, where $T_i \subseteq \mathcal{T}$ $\forall i$.*
- *$\mathbf{T} = (\mathbf{T}_1, \ldots, \mathbf{T}_n)$ is the vector of prior probability distributions. $\mathbf{T}_i(t_j)$ is the probability that player i is of type t_j. In this work we assume that all $\mathbf{T}_i$'s are independent.*
- *$A = A_1 \times \ldots \times A_n$ is the set of player actions, where $A_i \subseteq \mathcal{A}$ $\forall i$.*
- *$u = (u_1, \ldots, u_n)$ is the vector of payoff functions. For every i, the payoff of player i is $u_i : C_1 \times \ldots \times C_n \mapsto [0, 1]$, where $C_j = (T_j, A_j)$ is the set of possible type-action values for player j. We will also slightly abuse notation, and denote by $u_i(\mathbf{C}) = \mathrm{E}_{c \sim \mathbf{C}}[u_i(c)]$.*

A simultaneous-move Bayesian game is played as follows. Each player i fixes a strategy $\mathbf{A}_i$, which may be a function of her to-be-determined type $t \in T_i$. Nature assigns each player i a type t_i according to her corresponding prior probability distribution $\mathbf{T}_i$. Each player then chooses a move a_i according to the distribution $\mathbf{A}_i(t_i)$, the probability distribution over her actions given the chosen type. Every player then receives her payoff, $u_i((a_1, t_1), \ldots, (a_n, t_n))$.

We use standard notation as follows. For a subset $I \subseteq N$ and some vector $X = X_1 \times \ldots \times X_n$, let $X_I = \bigotimes_{i \in I} X_i$, and $X_{-I} = \bigotimes_{i \notin I} X_i$. Furthermore, $\mathbf{C}_{-i} : \mathbf{C}'_i = \mathbf{C}_1 \times \ldots \times \mathbf{C}_{i-1} \times \mathbf{C}'_i \times \mathbf{C}_{i+1} \times \ldots \times \mathbf{C}_n$. We also abuse notation, and write $\mathbf{C}_{-i} : \mathbf{A}'_i = \mathbf{C}_1 \times \ldots \times \mathbf{C}_{i-1} \times (\mathbf{A}'_i, \mathbf{T}_i) \times \mathbf{C}_{i+1} \times \ldots \times \mathbf{C}_n$ where $\mathbf{C}_i = (\mathbf{A}_i, \mathbf{T}_i)$. This is an abuse of notation because $\mathbf{A}'_i$ may be a function of $\mathbf{T}_i$.

In order to describe distributions $\mathbf{C}$ that correspond to the Nash equilibria of a particular game, we define Nash equilibrium type-action distributions.

Definition 2.2 (Nash equilibrium type-action distribution) *Let* $\mathbf{C} = \mathbf{C}_1 \times \ldots \times \mathbf{C}_n$, *where* $\mathbf{C}_i$ *is a distribution over* $T_i \times A_i$. $\mathbf{C}$ *is a* Nash equilibrium type-action distribution *if:*

- *The marginal distribution over players' types is* $(\mathbf{T}_1, \ldots, \mathbf{T}_n)$.
- *For every player* i, *type* $t \in T_i$, *and action* $a \in A_i$,

$$\mathop{\mathrm{E}}_{c \sim (\mathbf{C}|\mathbf{T}_i = t)} [u_i(c_{-i} : a)] \leq \mathop{\mathrm{E}}_{c \sim (\mathbf{C}|\mathbf{T}_i = t)} [u_i(c)] .$$

3 Main Results

Consider some Nash equilibrium type-action distribution $\mathbf{C}$ over C. This distribution is an ex post Nash equilibrium if for every player, her chosen strategy is a best-response even conditioned on the sampled values of other players. Formally,

Definition 3.1 (ex post Nash equilibrium) *A Nash equilibrium type-action distribution* $\mathbf{C}$ *is in an* ex post Nash equilibrium *if for every* $c = ((t_1, a_1), \ldots, (t_n, a_n)) \in \mathrm{supp}(\mathbf{C})$ *and for every* i, $u_i(c) \geq u_i(c_{-i} : (t_i, a_i'))$ *for any* $a_i' \in A_i$.

For a discussion of the strength of this property, its relation to purification, learning, and rational expectations, as well as its relevance to mechanism design and implementation theory, see Kalai (2004).

One way of viewing the property of ex post Nash is as follows: For every player, given her chosen type, her chosen strategy remains **best**-response even after she sees the realized types and chosen actions of **all** other players.

The above sentence has two bolded words, both of which we will weaken slightly. The weakened version of ex post Nash reads as follows: For every player, given her chosen type, her chosen strategy remains **almost** best-response **with high probability** even after she sees the realized types and chosen actions of **some** other players.

We note that Kalai (2004) also allows an almost best-response with high probability, but, in contrast, players may be exposed to the realized types and chosen actions of **all** other players. An additional difference is that Kalai (2004) shows that in uniformly equicontinuous games the ex post property holds with high probability for all players simultaneously, whereas we will show that in general games the weakened version of ex post Nash holds only for **most** players at once[1].

As mentioned in the introduction, one issue is what it means to see the types and actions of **some** other players. In order to formalize the manner in which exposed sets of players are chosen, we introduce the notion of an *exposure rule* $\mathbf{R}$. $\mathbf{R}$ is composed of n variables, $(\mathbf{R}_1, \ldots, \mathbf{R}_n)$, one per player. Each $\mathbf{R}_i$ specifies a distribution over subsets of the players, which will be the players to which i is exposed. More formally:

Definition 3.2 (exposure rules) $\mathbf{R} = (\mathbf{R}_1, \ldots, \mathbf{R}_n)$ *is an* exposure rule *if it is a joint distribution of* n *random variables* $\mathbf{R}_i$, *where* $\mathbf{R}_i$ *takes values in* $P(N \setminus \{i\})$ *(the powerset of* $N \setminus \{i\}$*).*

[1] In fact, Example 3.8 shows that in general it may be impossible to obtain the weakened version of ex post Nash for all players simultaneously.

We now define our weakened notion of ex post Nash. Essentially, the definition requires that for every player, with probability at least $1-p$, her expected payoff can not increase by more than α with a change of action, even after she is exposed to the values specified by the exposure rule $\mathbf{R}$. In order to further explain the definition, we contrast it with the definition of a Nash equilibrium type-action distribution $\mathbf{C}$ (see Definition 2.2). If $\mathbf{C}$ is such an equilibrium, then

$$\neg\exists a \in A_i \text{ s.t. } u_i(\mathbf{C}_{-i} : a|\mathbf{T}_i = t) - u_i(\mathbf{C}|\mathbf{T}_i = t) > 0.$$

That is, there is no better strategy than the equilibrium strategy, conditioned on the realized type of a player (since the strategy may depend on the type). We wish this to hold even when player i obtains more information than just her type, but also the realized types and chosen actions of the exposed players.

Consider some player i and some type-action distribution $\mathbf{C}$. Also consider a set $r_i \subseteq N \setminus \{i\}$ and some type-action profile c_{r_i} for the players in r_i. Then the distribution $(\mathbf{C}|\mathbf{C}_{r_i} = c_{r_i})$ is the conditional type-action distribution in which the players in r_i have type-action profile c_{r_i}. Ideally, we would like the following to hold:

$$\neg\exists a \in A_i \text{ s.t. } u_i(\mathbf{C}_{-i} : a|\mathbf{T}_i = t, \mathbf{C}_{r_i} = c_{r_i}) - u_i(\mathbf{C}|\mathbf{T}_i = t, \mathbf{C}_{r_i} = c_{r_i}) > 0.$$

This means that there is no better strategy than the equilibrium strategy, conditioned on the realized type of a player *and exposure to the set r_i when those players have type-action profile c_{r_i}*. However, this is asking too much, and so we weaken the requirement in two ways. First, we replace the 0 by some small parameter α, implying that i should not be able to improve her payoff by more than α. Second, instead of considering arbitrary r_i and c_{r_i}, we sample r_i from the distribution $\mathbf{R}_i$, and consider a random choice of the values of c_{r_i} sampled from $\mathbf{C}_{r_i}$. We then require that player i should not be able to increase her payoff by more than α with probability more than p over these random choices. The formal definition is the following.

Definition 3.3 ((p,α)-ex post Nash with exposure rule R) *A Nash equilibrium type-action distribution $\mathbf{C}$ is* (p,α)-ex post Nash with exposure rule $\mathbf{R}$ *if for every player i and every $t \in T_i$:*

$$\Pr_{\substack{c\sim(\mathbf{C}|\mathbf{T}_i=t)\\ r_i\sim\mathbf{R}_i}}[\exists a \in A_i \text{ s.t. } u_i(\mathbf{C}_{-i} : a|\mathbf{C}_{r_i} = c_{r_i}, \mathbf{T}_i = t) - u_i(\mathbf{C}|\mathbf{C}_{r_i} = c_{r_i}, \mathbf{T}_i = t) > \alpha] < p.$$

3.1 Random Exposure

Our first natural flavor of exposure rules is one in which exposed sets are chosen at random.

Definition 3.4 (randomized exposure rule randR^m) **randR**$^m = (\mathbf{R}_1, \ldots, \mathbf{R}_n)$, *where, for each i, $\mathbf{R}_i$ is the uniform distribution over all $R_i \subset N \setminus \{i\}$ of size $|R_i| = m$.*

We now state our main theorem on random exposure.

Theorem 3.5 *Fix any $p > 0$, $\alpha > 0$, and integer $m > 0$. Then in any game $G \in \Gamma(\mathcal{C})$, every Nash equilibrium type-action distribution $\mathbf{C}$ of G is*

$$\left(p + \frac{32m \cdot |\mathcal{A}|}{(n-m)p\alpha^2}, \alpha\right)\text{-ex post Nash with exposure rule } \mathbf{randR}^m.$$

A slightly different way of stating Theorem 3.5 is the following corollary, which follows immediately by rearranging the parameters:

Corollary 3.6 *For any $p > 0$, $\alpha > 0$, natural number n, and game $G \in \Gamma(\mathcal{C})$ with n players, every Nash equilibrium type-action distribution $\mathbf{C}$ of G is (p, α)-ex post Nash with exposure rule $\mathbf{randR}^m$ for any integer m satisfying*

$$m \leq \frac{p^2\alpha^2 \cdot (n-m)}{128 \cdot |\mathcal{A}|}.$$

Note that if p and α are bounded away from zero independently of n, then m can be as large as a constant fraction of all the players.

Alternatively, we can state the theorem as follows:

Corollary 3.7 *For any $p > 0$, $\alpha > 0$, and natural number m, every Nash equilibrium type-action distribution $\mathbf{C}$ of every game $G \in \Gamma(\mathcal{C})$ with at least n players is (p, α)-ex post Nash with exposure rule $\mathbf{randR}^m$, where*

$$n = \frac{128 \cdot |\mathcal{A}| \cdot m}{p^2\alpha^2} + m.$$

3.2 "Forbidden Set" Exposure

As mentioned in the introduction, the partial exposure ex post Nash property may not hold when players are exposed to arbitrary, as opposed to randomly chosen, players. Consider the following example, a game called Match-the-Expert:

Example 3.8 For every player i, u_i is a function of exactly one player, D_i, that outputs 1 on agreement and 0 otherwise. Consider the Nash equilibrium in which every player outputs 0 or 1 with equal probability.

Clearly, if any player i is exposed to her expert D_i, she will want to change her action with probability 1/2. However, note that in this case if player i is exposed to any player besides D_i, then her original strategy *is* still optimal. Essentially, our second flavor of exposure shows that this sort of decomposition between a few influential players and a majority of harmless ones occurs in every game.

"Forbidden set" exposure is roughly the following: for each player there will be a *small* set of other players B_i that she is forbidden from seeing. From the remaining players, however, she can choose an arbitrary subset of m players to see. To formalize this idea, we present the following definitions:

Definition 3.9 (m-bounded exposure rule) *An exposure rule $\mathbf{R} = (\mathbf{R}_1, \ldots, \mathbf{R}_n)$ is m-*bounded *if for each i and each $r_i \in \text{supp}(\mathbf{R}_i)$, the set $r_i \subset N \setminus \{i\}$ is of size $|r_i| \leq m$.*

Definition 3.10 (B-avoiding exposure rule) *Let $B = B_1 \times \ldots \times B_n$, where $B_i \subseteq N \setminus \{i\}$. Then an exposure rule $\mathbf{R} = (\mathbf{R}_1, \ldots, \mathbf{R}_n)$ is B-avoiding if for all i and all $r_i \in \text{supp}(\mathbf{R}_i)$ it holds that $|r_i \cap B_i| = \emptyset$.*

Our main theorem here will state that there exists some B such that if $\mathbf{R}$ is B-avoiding and m-bounded, then the Nash equilibrium type-action distribution is (p, α)-ex post Nash with exposure rule $\mathbf{R}$. Now, such a statement would be trivially true if we had $B_i = N \setminus \{i\}$ for all i, since this means that no player can be exposed to any other player (as all other players are forbidden to her). The whole point, then, is that the sets B_i will be **small**. There is clearly a tradeoff between m, p, α, and the size of each B_i, and our aim is to optimize this tradeoff. Ideally, we want to maximize m while minimizing p, α, and $|B_i|$. We note that the size of B_i also depends on the sizes of the action and type spaces, but, importantly, is independent of the number of players n.

We now state the main theorem. For this theorem, we restrict attention to a subclass of games $\Gamma(\mathcal{C}') \subset \Gamma(\mathcal{C})$, where $\mathcal{C}' \equiv \mathcal{A} \times \mathcal{T}'$ and $\mathcal{T}'$ is also a finite set (and not just $\mathcal{A}$). This is the class of all games in which both the type and action spaces are finite.

Theorem 3.11 *Fix some positive p, α, and integer m. Also fix a game $G \in \Gamma(\mathcal{C}')$, and let $\mathbf{C}$ be some Nash equilibrium type-action distribution of G. Then for every player i there exists a set $B_i \subseteq N \setminus \{i\}$ of size*

$$|B_i| \leq \frac{32\,|\mathcal{A}| \cdot |\mathcal{T}'| \cdot m}{p\alpha^2}$$

such that $\mathbf{C}$ is in a (p, α)-ex post Nash equilibrium with any m-bounded $(B_1, \ldots, B_n)$-avoiding exposure rule $\mathbf{R}$.

In general, we think of p and α as arbitrarily small constants independent of n. The sizes of $\mathcal{A}$ and $\mathcal{T}'$ are also finite, and are thus independent of n. In this case, if a player is exposed to the types and actions of m players of her choice (except from those in B_i), then the size of the forbidden set is $O(m)$ – that is, there exists some universal constant γ such that $|B_i| < \gamma m$ for all m and all games in $\Gamma(\mathcal{C}')$.

Note that if we wish to allow each player to be exposed to a constant fraction of players (i.e. if we want to have m/n be a constant), we can do so by forbidding her access to some (larger) constant fraction (namely, $m/np\alpha^2$) of players.

It is not necessary to choose p and α independently of n. However, if p and α are chosen to be too small (relative to n), then the statement of Theorem 3.11 will be trivially true, since the guaranteed forbidden set will be larger than the total number of players. This is unavoidable – the example in Section 5.1 shows that, in general, it is necessary to have at least $m/p\alpha^2$ forbidden players.

Larger type space for "forbidden set" exposure One of the differences between random exposure and the forbidden set exposure is that the latter is meaningful only when the size of the type space is finite (since the upper bound on $|B_i|$ contains the term $|\mathcal{T}|$). It is possible to get around this if we allow a different forbidden set for every type of player i. That is, for every $t \in T_i$, the forbidden set is some B_i^t. The order of play is now as follows: first Nature chooses a type t_i for each player i, and each player chooses an action according

to the corresponding strategy. Next, each player i is exposed according to a $(B_1^{t_1}, \ldots, B_n^{t_n})$-avoiding exposure rule $\mathbf{R}^m$. Then a similar statement to that of Theorem 3.11 would hold: with high probability, a player will gain little by a change of action.

3.3 Stability Under Costly Exposure

As noted in the introduction, an interesting interpretation of our result is that, when given an opportunity to be exposed to the realized types and chosen actions of other players and to change their own strategy (at a small cost), players will decline. Thus, the original equilibria will survive. This is formalized in the following theorem.

Theorem 3.12 *Consider any game $G \in \Gamma(\mathcal{C})$, and let G' be the game with the following modification:*

- *All players' types are realized, and players choose a strategy to play according to some Nash equilibrium type-action distribution $\mathbf{C}$.*
- *Each player is then given an opportunity to be exposed to the realized types and chosen actions of other players, according to an exposure rule $\mathbf{R}$. The only condition is that $\mathbf{C}$ be (p, α)-ex post Nash with exposure rule $\mathbf{R}$ (for example, random exposure or B-avoiding exposure). Following this exposure, she may change her strategy given the new information. The cost of the exposure and revision is at least $p + \alpha$.*

Then if all players play according to $\mathbf{C}$, no rational player will take this opportunity. That is, the original equilibria of G will survive.

Suppose now that players are also allowed to choose the number of players to whom they are randomly exposed. Suppose further that there is some cost c per player for this exposure – i.e., a player who wishes to be randomly exposed to the realized types and chosen actions of m players must pay $c \cdot m$. Then the following theorem shows that if c is set correctly, players will again decline to be exposed (to any player), and the original Nash equilibria will survive. Note that c below depends only on $|\mathcal{A}|$.

Theorem 3.13 *Consider any game $G \in \Gamma(\mathcal{C})$, and let G' be the game with the following modification:*

- *All players' types are realized, and players choose a strategy to play according to some Nash equilibrium type-action distribution $\mathbf{C}$.*
- *Each player then picks some natural number m, and is given the opportunity to be exposed to the realized types and chosen actions of m random players players (i.e., according to exposure rule $\mathbf{randR}^m$). The cost of the exposure is c for each player, or $c \cdot m$, where*

$$c \geq 3 \cdot \sqrt[4]{32\,|\mathcal{A}|}.$$

Then if all players play according to $\mathbf{C}$, no rational player will take this opportunity. Again, the original equilibria of G will survive.

4 Proofs

In this section we will prove our main theorems. We will argue that when a player is exposed to the realized types and chosen actions of other players according to some exposure rule, then the values she sees do not affect her payoff too much. Intuitively, this will hold if the players whose values she saw do not have too much influence on her payoff. Thus, we begin with a notion of the influence of players in functions (Section 4.1), and then proceed to relate it to exposure and robustness to exposure in games (Section 4.2).

4.1 Pivotal Players and Sets

The notion of influence that we need was introduced by Al-Najjar and Smorodinsky (2000) and further studied by Haggstrom et al. (2006) and Gradwohl et al. (2008). We note that similar notions were studied by Malaith and Postlewaite (1990) and by Fudenberg et al. (1998).

Let $F : C_1 \times \ldots \times C_n \mapsto [0,1]$ be some function, where each player provides her corresponding input. Suppose each player i's type-action character is drawn from a distribution $\mathbf{C}_i$ over C_i, and denote by $\mathbf{C}$ the product distribution of the $\mathbf{C}_i$'s. The notion of influence we need is the following:

Definition 4.1 (pivotal player) *Player i is (p,α)-pivotal in F with respect to $\mathbf{C}$ if*

$$\Pr_{c_i \sim \mathbf{C}_i}\left[\left|\,\mathrm{E}[F] - \mathrm{E}[F|\mathbf{C}_i = c_i]\right| \geq \alpha\right] \geq p.$$

Denote the number of (p,α)-pivotal players in F with respect to $\mathbf{C}$ by $K(F,\mathbf{C},p,\alpha)$. The following theorem is a restriction of a theorem of Al-Najjar and Smorodinsky (2000) to the case of a countable support. This formulation of the theorem is given by Gradwohl et al. (2008).

Lemma 4.2 (Al-Najjar and Smorodinsky 2000, Theorem 4) *Let $\mathbf{C} = \mathbf{C}_1 \times \ldots \times \mathbf{C}_n$ be some distribution over $\mathcal{C}^n$. Then for any $0 < \alpha < 1$, $0 < p < 1$, and $F : C_1 \times \ldots \times C_n \mapsto [0,1]$,*

$$K(F,\mathbf{C},p,\alpha) \leq \frac{8}{p\alpha^2}.$$

A definition similar to that of pivotal players also holds for sets of players:

Definition 4.3 (pivotal set of players) *A set of players R is (p,α)-pivotal in F with respect to $\mathbf{C}$ if*

$$\Pr_{c_R \sim \mathbf{C}_R}\left[\left|\,\mathrm{E}[F] - \mathrm{E}[F|\mathbf{C}_R = c_R]\right| \geq \alpha\right] \geq p.$$

In an arbitrary function F, there may be many sets of players with large influence. Suppose, for example, that there is some player with large influence. Then any set that contains her will also have large influence. We would like to argue that there are few players that, when eliminated, will leave the remaining players with little influence **even in sets**. To this end, we define the exclusively-influential set – a set of players that, when removed, leaves all subsets of the remaining players with little influence.

Definition 4.4 ((m,p,α)-exclusively-influential set for $(F,\mathbf{C})$) *A set $S \subset N$ is an (m,p,α)-*exclusively-influential set for $(F,\mathbf{C})$ *if $\forall R \subset N \setminus S$ of size $|R| \leq m$, R is **not** (p,α)-pivotal in F with respect to $\mathbf{C}$.*

Lemma 4.5 shows that there always exists an exclusively-influential set of size at most $2m/p\alpha^2$.

Lemma 4.5 *Fix $0 < \alpha < 1$, $0 < p < 1$, and let $\mathbf{C}$ be as in Lemma 4.2. Then for any F and natural number m, $\exists S \subset N$ of size $|S| \leq \frac{8m}{p\alpha^2}$, such that $\forall R \subset N \setminus S$, if $|R| \leq m$ then R is **not** (p,α)-pivotal.*

Proof: We first bound the maximal number of *disjoint* influential sets.

For $i = 1, \ldots, k$, define subsets $S_i \subset N$ of size $|S_i| \leq m$ satisfying the following:

- All S_i's are disjoint.
- For all i, S_i is (p,α)-pivotal.

Let $\{S_i\}_{i=1}^k$ be some maximal collection of such sets (i.e. any other (p,α)-pivotal set of size at most m intersects one of the S_i's.). We will bound k.

To simplify the exposition, suppose all S_i's are of size m, S_1 is comprised of players $1, \ldots, m$, S_2 is comprised of players $m+1, \ldots, 2m$, and so on. Now consider the function $F' : \mathbf{C}'_1 \times \ldots \times \mathbf{C}'_t \times \mathbf{C}_{mt+1} \times \ldots \times \mathbf{C}_n \mapsto [0,1]$, where $\mathbf{C}'_i = \mathbf{C}_{im+1} \times \ldots \times \mathbf{C}_{(i+1)m}$. This function takes the same values as F, except that it considers the inputs of all the players in S_i as the input of one meta-player. Clearly, F' has the same expectation as F.

By Lemma 4.2, the number of (p,α)-pivotal meta-players in F' is at most $8/p\alpha^2$. Note that since we assumed that the set of players S_i is (p,α)-pivotal in F, we know that the meta-player representing the set S_i is (p,α)-pivotal in F'. Thus, we can conclude that $k \leq 8/p\alpha^2$.

Let $S = \bigcup_i S_i$. Now consider some set $R \subset N$ of size $|R| \leq m$. If R is disjoint from the S_i's, then R can **not** be (p,α)-pivotal, since the S_i's are a maximal collection of such disjoint influential sets. Conversely, if R is (p,α)-pivotal, then $R \cap S \neq \emptyset$.

We are now done: $|S| \leq 8m/p\alpha^2$ as required, and for any $R \subset N \setminus S$ of size $|R| \leq m$, we have that R is **not** (p,α)-pivotal. ■

4.2 Proofs of Partial-Exposure Ex Post Nash

In this section we prove Theorem 3.5 on random exposure and Theorem 3.11 on "forbidden set" exposure. We begin with Lemma 4.6, which will be used in both proofs.

Lemma 4.6 *Let $G \in \Gamma(\mathcal{C})$, and fix some Nash equilibrium type-action distribution $\mathbf{C}$ of G and exposure rule $\mathbf{R} = (\mathbf{R}_1, \ldots, \mathbf{R}_n)$. Suppose that for all players i, with probability $1-q$ (over the choice of $r_i \sim \mathbf{R}_i$ and $c \sim (\mathbf{C}|\mathbf{T}_i = t)$) it holds that*

$$|u_i(\mathbf{C}_{-i} : a|\mathbf{C}_{r_i} = c_{r_i}, \mathbf{T}_i = t) - u_i(\mathbf{C}_{-i} : a|\mathbf{T}_i = t)| \leq \alpha$$

for all $a \in A_i$ and $t \in T_i$. Then $\mathbf{C}$ is $(q, 2\alpha)$-ex post Nash with exposure rule $\mathbf{R}$.

The intuition behind Lemma 4.6 is the following: for "good" c and r_i, i.e. ones that satisfy the inequality in the statement of the lemma, exposure does not change the expected utility of the player by much, regardless of her action. In particular, exposure does not decrease her expected utility when she plays her Nash strategy, nor does it increase the expected utility when she plays an alternate strategy. Hence, the Nash strategy is close to optimal.

Proof: Fix some $t \in T_i$. Call c and r_i "good" if for all $a \in A_i$

$$|u_i(\mathbf{C}_{-i} : a|\mathbf{C}_{r_i} = c_{r_i}, \mathbf{T}_i = t) - u_i(\mathbf{C}_{-i} : a|\mathbf{T}_i = t)| \leq \alpha.$$

Note that $\Pr[c \text{ and } r_i \text{ are "good"}] \geq 1 - q$. For any good c and r_i, let $\mathbf{A}'_i = \mathbf{A}'_i(c_{r_i}, t)$ be some fixed strategy. Player i can play this strategy regardless of the value of c. Thus

$$\begin{aligned}
&\left|u_i(\mathbf{C}_{-i} : \mathbf{A}'_i|\mathbf{C}_{r_i} = c_{r_i}, \mathbf{T}_i = t) - u_i(\mathbf{C}_{-i} : \mathbf{A}'_i|\mathbf{T}_i = t)\right| \\
&= \left|\sum_{a \in A_i} \Big(u_i(\mathbf{C}_{-i} : a|\mathbf{C}_{r_i} = c_{r_i}, \mathbf{T}_i = t) - u_i(\mathbf{C}_{-i} : a|\mathbf{T}_i = t)\Big) \cdot \Pr\left[\mathbf{A}'_i = a\right]\right| \\
&\leq \sum_{a \in A_i} \alpha \cdot \Pr\left[\mathbf{A}'_i = a\right] = \alpha. \qquad (1)
\end{aligned}$$

Now, for good c and r_i, we get that for any action $a \in A_i$,

$$\begin{aligned}
&u_i(\mathbf{C}_{-i} : a|\mathbf{C}_{r_i} = c_{r_i}, \mathbf{T}_i = t) \\
&\qquad \leq u_i(\mathbf{C}_{-i} : a|\mathbf{T}_i = t) + \alpha \\
&\qquad \leq u_i(\mathbf{C}|\mathbf{T}_i = t) + \alpha \\
&\qquad \leq u_i(\mathbf{C}|\mathbf{C}_{r_i} = c_{r_i}, \mathbf{T}_i = t) + 2\alpha,
\end{aligned}$$

where the first and third inequalities follow from (1), first with $\mathbf{A}'_i = a$ and then with $\mathbf{A}'_i = \mathbf{A}_i$. The second inequality follows from the fact that $\mathbf{C}$ is a Nash equilibrium type-action distribution.

Since good c and r_i are chosen with probability at least $1 - q$, the claim follows. ■

4.2.1 Proof of Random Exposure Theorem

We now prove Theorem 3.5.

Proof: Fix some game $G \in \Gamma(\mathcal{C})$, a Nash equilibrium type-action distribution $\mathbf{C}$ of G, a player i and a type $t \in T_i$. Also fix some positive p, α, and m.

Recall that the random exposure rule $\mathbf{randR}^m = (\mathbf{R}_1, \ldots, \mathbf{R}_n)$ exposes player i to the realized types and chosen actions of a random subset of m players. We need to show that for every $a \in A_i$,

$$\Pr_{\substack{r_i \sim \mathbf{R}_i \\ c \sim (\mathbf{C}|\mathbf{T}_i = t)}} [u_i(\mathbf{C}_{-i} : a|\mathbf{C}_{r_i} = c_{r_i}, \mathbf{T}_i = t) - u_i(\mathbf{C}|\mathbf{C}_{r_i} = c_{r_i}, \mathbf{T}_i = t) > 2\alpha] < q$$

for a suitable choice of q.

One way of choosing m random players from $N \setminus \{i\}$ is via the following procedure:

(1.) Partition the $n-1$ players of $N \setminus \{i\}$ into sets of size m uniformly at random (the last set might be of size less than m).

(2.) Choose one of the sets of size m uniformly at random (ignoring the last set if it is smaller).

Step (1.) of the procedure partitions $N \setminus \{i\}$ into $\lceil n/m \rceil$ sets of players. We can view each such set as a meta-player – a player whose output is the output of the m players in the set. Consider some partitioning of the players to sets $S_1, \ldots, S_k$, where S_k is the set or meta-player that may consist of fewer than m players. Now consider the function $F^a : C_{S_1} \times \ldots \times C_{S_k} \mapsto [0,1]$ defined as follows:

- For each $i \in \{1, \ldots, k\}$, let s_i be some element of
$$C_{S_i} = \bigotimes_{j \in S_i} C_j.$$
- Then $F^a(s_1, \ldots, s_k) = u_i(c_1, \ldots, c_n)$, the utility function of player i given type-action pairs $c_1, \ldots, c_n$, where
  - $c_i = (t, a)$, and
  - for all $j \neq i$ it holds that $c_j = (t_j, a_j)$, where (t_j, a_j) is the type-action pair specified in s_ℓ for player j by the meta-player S_ℓ such that $j \in S_\ell$.

By Lemma 4.2, the number of (p,α)-pivotal players in F^a is at most $8/p\alpha^2$. Note that if the meta-player S_j is **not** (p,α)-pivotal in F^a, then

$$\left|u_i(\mathbf{C}_{-i} : a | \mathbf{C}_{S_j} = c_{S_j}, \mathbf{T}_i = t) - u_i(\mathbf{C}_{-i} : a | \mathbf{T}_i = t)\right| \leq \alpha$$

with probability at least $1-p$ (over the choice of $c \sim (\mathbf{C}|\mathbf{T}_i = t)$).

Now, in step (2.) of the procedure we choose one of $S_1, \ldots, S_{k-1}$ at random. Since there are at most $8/p\alpha^2$ meta-players that are (p,α)-pivotal in F^a, the probability of choosing one is at most

$$q = \frac{8}{p\alpha^2 \cdot \lfloor n/m \rfloor} \leq \frac{8m}{(n-m)p\alpha^2}.$$

The probability that there exists any $a \in A_i$ such that the chosen meta-player is (p,α)-pivotal in F^a is at most $|A_i| \cdot q$.

Putting the above together, we get that the probability that $\exists a \in A_i$ such that

$$\left|u_i(\mathbf{C}_{-i} : a | \mathbf{C}_{S_j} = c_{S_j}, \mathbf{T}_i = t) - u_i(\mathbf{C} : a | \mathbf{T}_i = t)\right| > \alpha$$

is at most $p + |A_i| \cdot q$ over the choice of $c \sim (\mathbf{C}|\mathbf{T}_i = t)$ and the uniform choice of S_j from all sets of size m. The $|A_i| \cdot q$ term comes from the probability that S_j is (p,α)-pivotal for some $a \in A_i$, and the p term comes from the probability of changing the utility when S_j is not (p,α)-pivotal.

What we have shown thus far is that with probability $1 - p - |A_i| \cdot q$ (over the choice of $r_i \sim \mathbf{R}_i$ and $c \sim (\mathbf{C}|\mathbf{T}_i = t)$) it holds that for all $a \in A_i$ and $t \in T_i$,

$$|u_i(\mathbf{C}_{-i} : a | \mathbf{C}_{r_i} = c_{r_i}, \mathbf{T}_i = t) - u_i(\mathbf{C}_{-i} : a | \mathbf{T}_i = t)| \leq \alpha.$$

Applying Lemma 4.6 completes the proof. ■

4.2.2 Proof of "Forbidden Set" Exposure Theorem

In this section we prove Theorem 3.11.

Proof: For each i, each $t \in T_i$, and each $a \in A_i$, let $S^{t,a}$ be an (m, p, α)-exclusively-influential set for $(u_i, (\mathbf{C}|\mathbf{A}_i = a, \mathbf{T}_i = t))$ guaranteed by Lemma 4.5. Let

$$B_i = \bigcup_{\substack{a \in A_i \\ t \in T_i}} S^{t,a}.$$

By Lemma 4.5,

$$\left|S^{t,a}\right| \leq \frac{8m}{p\alpha^2},$$

and so

$$|B_i| \leq \frac{8\,|\mathcal{A}| \cdot |\mathcal{T}'| \cdot m}{p\alpha^2}$$

as claimed.

Now consider some m-bounded $(B_1, \ldots, B_n)$-avoiding exposure rule $\mathbf{R} = (\mathbf{R}_1, \ldots, \mathbf{R}_n)$, and fix some player i. We wish to show that $\mathbf{C}$ is (p, α)-ex post Nash with exposure rule $\mathbf{R}$. By the definition of B_i, $\mathbf{R}_i$ has the following property:

For every $r_i \in \text{supp}(\mathbf{R}_i)$ and $a \in A_i$, r_i is **not** (p, α)-pivotal for $(u_i, (\mathbf{C}|\mathbf{A}_i = a, \mathbf{T}_i = t))$. Note that $u_i(\mathbf{C}|\mathbf{A}_i = a, \mathbf{T}_i = t) = u_i(\mathbf{C}_{-i} : a|\mathbf{T}_i = t)$ (this is just a difference of notation). Thus,

$$\Pr_{c \sim (\mathbf{C}|\mathbf{T}_i = t)} [|u_i(\mathbf{C}_{-i} : a|\mathbf{C}_{r_i} = c_{r_i}, \mathbf{T}_i = t) - u_i(\mathbf{C}_{-i} : a|\mathbf{T}_i = t)| > \alpha] < p.$$

Since this holds for all $r_i \in \text{supp}(\mathbf{R}_i)$, it holds in particular when we pick r_i at random from some distribution $\mathbf{R}_i$. Hence,

$$\Pr_{\substack{c \sim (\mathbf{C}|\mathbf{T}_i = t) \\ r_i \sim \mathbf{R}_i}} [|u_i(\mathbf{C}_{-i} : a|\mathbf{C}_{\mathbf{R}_i} = c_{r_i}, \mathbf{T}_i = t) - u_i(\mathbf{C}_{-i} : a|\mathbf{T}_i = t)| > \alpha] < p.$$

For every $t \in T_i$, the probability that for all $a \in A_i$ simultaneously we have

$$|u_i(\mathbf{C}_{-i} : a|\mathbf{C}_{\mathbf{R}_i} = c_{r_i}, \mathbf{T}_i = t) - u_i(\mathbf{C}_{-i} : a|\mathbf{T}_i = t)| \leq \alpha$$

is at least $1 - |A_i| \cdot p$. Applying Lemma 4.6 completes the proof. ■

4.3 Proof of Stability Under Costly Exposure Theorems

In this section we prove Theorems 3.12 and 3.13. We begin with the former.

Proof of Theorem 3.12: Suppose all players play according to a Nash equilibrium. Then the expected payoff of a player i is some value Q. We now compute the expected payoff if she takes the opportunity to be exposed, and show that it is strictly less than Q.

If player i is exposed according to an exposure rule $\mathbf{R}$, then the probability that she can improve her payoff by more than α is at most p. With probability p, she can improve her payoff to at most 1 (since the payoff is bounded). Thus, her payoff is

$$Q' \leq p + (1 - p)(Q + \alpha) - \alpha - p < Q.$$

This implies that the preferred strategy is to not be exposed. ■

Proof of Theorem 3.13: Theorem 3.5 states that for any $p > 0$, $\alpha > 0$, and integer $m > 0$, $\mathbf{C}$ is

$$\left(p + \frac{32m \cdot |\mathcal{A}|}{(n-m)p\alpha^2}, \alpha\right)\text{-ex post Nash with exposure rule } \mathbf{randR}^m.$$

Fix $p = \alpha = \sqrt[4]{32\,|\mathcal{A}|}$. Plugging in these values we get that, for any $m > 0$, $\mathbf{C}$ is (p', α')-ex post Nash with exposure rule $\mathbf{randR}^m$, where

$$p' = \sqrt[4]{32\,|\mathcal{A}|} + \frac{m}{n-m} \cdot \sqrt[4]{32\,|\mathcal{A}|} \leq (m+1) \cdot \sqrt[4]{32\,|\mathcal{A}|}$$

and

$$\alpha' = \sqrt[4]{32\,|\mathcal{A}|}.$$

Now, a player who is exposed to m players must pay $m \cdot c \geq 3m \cdot \sqrt[4]{32\,|\mathcal{A}|} \geq p' + \alpha'$. We can continue exactly as in Theorem 3.12: Players are exposed via an exposure rule that is (p', α')-ex post Nash, but the cost of this exposure is at least $p' + \alpha'$. The argument of Theorem 3.12 shows that it is a better response to not be exposed. ■

5 Lower Bounds

Our second main result, Theorem 3.11, takes parameters m, p, and α, and shows that if a player is "forbidden" some set of players and is then exposed to m others of her choice, then with probability p she can not improve her payoff by more than α. A couple of natural questions to ask regarding the "tightness" of this characterization are: how many players must be forbidden, and how small can p be made? In this section we give two examples of games. The first shows that many players must be forbidden, since otherwise an equilibrium distribution of this game will not be (p, α)-ex post Nash. The second example shows that if the exposed sets of size βn are chosen at random, then there exists a game in which $p \geq \beta$. In particular, since $\beta \geq 1/n$, this means that in expectation at least one player *will* be able to improve her payoff by 1/2.

5.1 Lower Bound for Forbidding Players

The Match-the-Expert game of Example 3.8 is a game in which at least one player must be forbidden in order to have PE ex post Nash. In this section we show that sometimes even more players must be forbidden. We will show that for every choice of $p > 0$ and m there exists a game in which roughly the following holds: there is an $\alpha > 0$ such that if player i's forbidden set B_i is smaller than $m/p\alpha^2$, then she will be able to improve her payoff by more than α with probability more than p after exposure to set of m players disjoint from B_i. In other words, the equilibrium will not be (p, α)-ex post Nash.

The game is a variation on a game from Al-Najjar and Smorodinsky (2000). Players $j \neq i$ have action space $\{0, 1, \perp\}$, and we consider some Nash equilibrium in which they all output 0 or 1 with probability $\gamma/2$ (γ will be determined later), and $\perp$ with probability $1 - \gamma$. Player i outputs 0 or 1, each with probability 1/2. There is some set $I \subseteq N \setminus \{i\}$ of "important" players, and player i wins if and only she matches the majority of the

important players who did not output $\perp$. With high probability, the number of players who do not output $\perp$ is roughly $\gamma \cdot |I|$. Now, suppose i observes the actions of some set $S \subset I$ of size $|S| = m$. The probability that some player in S did not output $\perp$ is roughly γm. Conditioned on this happening, the expectation of the majority is now shifted (either towards 0 or towards 1, depending on what the one player played) by roughly $1/\sqrt{\gamma \cdot |I|}$. Thus, with probability γm, player i can improve her expected payoff by $1/\sqrt{\gamma \cdot |I|}$. That is, the equilibrium is not (p, α)-ex post Nash for $p = \gamma m$ and $\alpha = 1/\sqrt{\gamma \cdot |I|}$.[2] In order to maintain the ex post Nash property, player i must be forbidden from observing *any* set of m important players. That is, the forbidden set must be of size at least $|I| - m + 1$. If $|I|$ is much larger than m, this is roughly equal to $|I|$, which in turn is equal to $m/p\alpha^2$.

Finally, we note that this example can be slightly modified to show that the dependence on $|\mathcal{A}|$ and $|\mathcal{T}|$ is also necessary.

5.2 Lower Bound for p (Ex Post Only for Most Players)

In this section we define a game for which the Nash equilibrium is ex post for most players, but not all, when each player is exposed to a random set of size βn. The game is Match-the-Expert of Example 3.8.

If i is exposed to βn random players, the probability that she "sees" her expert D_i is β, in which case she will change her strategy and improve her payoff by 1/2. Thus, with exposure rule $\mathbf{randR}^{\beta n}$ and for any $\alpha \leq 1/2$ we have that $p \geq \beta$. If each player is exposed to at least one other player, and so $\beta \geq 1/n$, then the expected number of players who *will* want to change their actions is $pn \geq 1$. Thus, we can not expect to have the ex post Nash property hold for all players simultaneously.

References

[1] Y. Azrieli. Categorizing others in a large game. Games and Economic Behavior, forthcoming.

[2] N. I. Al-Najjar and R. Smorodinsky. Pivotal players and the characterization of influence. Journal of Economic Theory 92 (2000): 318–342.

[3] M. Blonski. Characterization of pure strategy equilibria in finite anonymous games. Journal of Mathematical Economics 34 (2000): 225–233.

[4] G. Carmona and K. Podczeck. On the existence of pure strategy Nash equilibria in large games. Journal of Economic Theory 144 (2009): 1300-1319.

[5] G. Carmona. Purification of Bayesian-Nash equilibria in large games with compact type and action spaces. Journal of Mathematical Economics 44 (2008): 1302–1311.

[6] E. Cartwright and M. Wooders. On equilibrium in pure strategies in games with many players. International Journal of Game Theory 38 (2009): 137–153.

[2] Note that since we can choose γ, we have control over the value of p. Since we can also choose $|I|$ we actually have much control over α as well, just subject to the constraints that $|I|$ must be an integer much larger than m.

[7] D. Fudenberg, D. Levine, and W. Pesendorfer. When are non-anonymous players negligible? Journal of Economic Theory 79 (1998): 46–71.

[8] R. Gradwohl and O. Reingold. Correlated types in large games. In preparation.

[9] R. Gradwohl and O. Reingold. Fault Tolerance in Large Games. Proceedings of the Ninth ACM Conference on Electronic Commerce, EC 2008.

[10] R. Gradwohl, O. Reingold, A. Yadin and A. Yehudayoff. Players' effects under limited independence. arxiv:math.PR:0805.0400.

[11] O. Haggstrom, G. Kalai and E. Mossel. A law of large numbers for weighted majority. Advances in Applied Mathematics 37:1 (2006): 112–123.

[12] J. Halpern. A computer scientist looks at game theory. Games and Economic Behavior 45:1 (2003): 114–131.

[13] W. Hoeffding. Probability inequalities for sums of bounded random variables. Journal of the American Statistical Association 58 (1963): 13–30.

[14] E. Kalai. Large robust games. Econometrica 72:6 (2004): 1631–1665.

[15] E. Kalai. Partially-specified large games. Lecture Notes in Computer Science 3828 (2005): 3–13.

[16] G. J. Mailath and A. Postlewaite. Asymmetric information bargaining problems with many agents. Review of Economic Studies 57 (1990): 351–167.

[17] D. Schmeidler. Equilibrium points of nonatomic games. Journal of Statistical Physics 17 (1973): 295–300.

[18] M. Wooders, E. Cartwright, and R. Selten. Behavioral conformity in games with many players. Games and Economic Behavior 57 (2006): 347–360.