

Robustness of Equilibria in Anonymous Local Games*

Willemien Kets[†]

October 12, 2010

Abstract

This paper studies the robustness of symmetric equilibria in anonymous local games to perturbations of prior beliefs. Two priors are strategically close on a class of games if players receive similar expected payoffs in equilibrium under the priors, for any game in that class. I show that if the structure of payoff interdependencies is sparse in a well-defined sense, the conditions for strategic proximity in anonymous local games are strictly weaker than the conditions for general Bayesian games of Kajii and Morris [11] when attention is restricted to symmetric equilibria. Hence, by exploiting the properties of anonymous local games, it is possible to obtain stronger robustness results for this class.

JEL classification: C72, D82

Keywords: Incomplete information, higher-order beliefs, robustness, strategic proximity, local interactions, anonymous games, heterogeneous externalities, network games.

*This paper subsumes the papers “Convergence of beliefs in Bayesian network games” and “Beliefs in Network Games,” which form Chapters 3 and 4 of my dissertation, respectively. I am indebted to Carlos Alós-Ferrer, Toni Calvó-Armengol, Sanjeev Goyal, Wouter Kager, John Miller, Stephen Morris, Jakub Steiner, Dolf Talman, Eric van Damme, Fernando Vega-Redondo, Mark Voorneveld, and, especially, Matthew Jackson for helpful comments and suggestions. I am grateful to the referees and the Editor, Christian Hellwig, for excellent comments that significantly improved the paper. All remaining errors are, of course, my own.

[†]Santa Fe Institute and CentER, Tilburg University. Address: 1399 Hyde Park Road, Santa Fe, NM, 87501. E-mail: willemien.kets@santafe.edu. Phone: 1-505-204-8012. Fax: 1-505-982-0565.

1 Introduction

As is well known, game-theoretic predictions can be highly sensitive to assumptions about players' higher-order beliefs (e.g., [3, 23]). This important insight has motivated the study of conditions under which predictions are robust to various perturbations of players' higher-order beliefs.¹ Unfortunately, the class of perturbations under which predictions are robust in all possible games is only small, underscoring the sensitivity of predictions. If, however, outcomes are robust to a wider range of perturbations in a subclass of games, then we need to be less concerned with the potential sensitivity of predictions when that class of games models the strategic situation we are interested in. This paper therefore studies the robustness of equilibria in a subclass of games that models a wide range of economic settings. It is shown that while higher-order beliefs can have a nontrivial effect on outcomes in these games, stronger robustness results can nevertheless be obtained for this class of games under certain conditions.

More specifically, one of the approaches to studying the robustness of strategic outcomes to small changes in players' beliefs is to identify a measure on information systems such that if two information systems are close according to this measure, then the predictions for any given Bayesian game are similar under the two information systems [11, 16]. In that case, the information systems are said to be *strategically close* on the class of all Bayesian games. If two information systems are strategically close in this sense, a modeler who is unsure which of the two information systems obtains can be confident that for any Bayesian game he may want to consider, the two information systems give similar predictions for that game.

However, if the information systems of interest do not meet this condition, the modeler's predictions are not necessarily sensitive to the choice of information system, as the information systems may give similar predictions for the subclass of games he is interested in. It is therefore important to identify classes of games of economic interest in which predictions are more robust to the choice of information system. In particular, since many economic settings are characterized by specific types of payoff interdependencies or symmetry properties, it is worth exploring whether such properties can be exploited to weaken the conditions for strategic proximity on a class of games that satisfy these properties.

This paper considers the strategic proximity of players' prior beliefs on the class of anonymous local games, a class of games that includes games with type-dependent heterogeneous externalities [2, 9] and network games [7]. In an anonymous local game, the payoffs to a player only depend on the actions and types of some subset of players, and do so in an anonymous way, in the sense that payoff functions do not depend on the identity of players. Moreover, each player faces the same environment before learning his type. Given the symmetry of the situation, it is natural to focus on symmetric equilibria.

¹See, among others, [4, 6, 10, 11, 16, 24, 25].

Instead of requiring information systems to give similar equilibrium predictions in every Bayesian game, as in much of the literature, I only require robustness of the symmetric equilibria in anonymous local games. More precisely, two (common) priors μ, μ' are said to be strategically close on the class of anonymous local games if, for any anonymous local game with prior μ and a symmetric (interim) equilibrium σ , there exists an approximate symmetric equilibrium σ' in the associated game with prior μ' such that the (ex ante) payoffs under σ and σ' are close.

Theorem 4.1 shows that two priors are strategically close on the class of anonymous local games if and only if (i) the priors assign similar ex ante probability to events; and (ii) with high prior probability, a player's type belongs to a cohesive subset of the set of types that have similar conditional beliefs under the two priors. This is the set of types with conditional beliefs that are close under the two priors such that with high conditional probability, their payoffs depend only on types with close conditional beliefs such that with high conditional probability, . . . their payoffs depend only on types with close conditional beliefs, for any number of iterations. Since these conditions make references to types only, and not to players, they characterize strategic proximity even if the player set is potentially unbounded (under the mild assumption that the expected number of players is finite).

Because strategic proximity is required to hold only for a subclass of games, the conditions for strategic proximity for general Bayesian games must be at least as strong as the present conditions if attention is restricted to symmetric equilibria. I therefore investigate the conditions under which the restriction to anonymous local games leads to conditions for strategic proximity that are *strictly* weaker than the conditions obtained by Kajii and Morris [11] for all Bayesian games. In particular, Kajii and Morris [11] show that a necessary condition for strategic proximity of priors on the class of Bayesian games is that with high prior probability, there is approximate common belief that conditional beliefs of all players are close under the priors.

Section 4.2 shows that the present conditions for strategic proximity are weaker than those for general Bayesian games if the following holds: There is a sequence of players i_1, \dots, i_k such that with high prior probability, player i_1 believes (with high conditional probability) that player i_2 believes that . . . the conditional beliefs of player i_k are not close under the two priors, and, in each such sequence, there is a player i_ℓ who has a vanishingly small effect on the payoffs of player $i_{\ell+1}$. The intuition is simple. If there is a player who believes that some other player believes . . . there is a player with conditional beliefs are not close under the two priors, then there is no approximate common belief of close conditional beliefs, so that the condition for general Bayesian games is not met. However, in anonymous local games, the (higher-order) beliefs of a player about the actions of certain other players may not affect his best response, so that predictions for this class of games will be similar under the two priors. This result may

be especially useful when the modeler is interested in large games, since the assumptions that interactions are local and anonymous are especially natural in such settings, while the conditions for general Bayesian games (that require approximate common belief) appear most likely to be violated.

The question which topologies on information structures induce continuity of the set of solutions has received considerable attention in the literature.² A common theme in this literature is that such a topology needs to be fairly strong, or, equivalently, that the class of perturbations to which outcomes are robust is small. I therefore depart from much of the literature by restricting attention to a specific subclass of Bayesian games, to investigate the extent to which the conditions on perturbations can be relaxed for that class.³

Of course, an important question is how to trade off the restrictions on the class of games with the extent to which continuity conditions can be relaxed. In the extreme, one could restrict attention to a single game, and get very weak conditions for strategic proximity for that game. It is therefore important to focus on a class of games such that the conditions for strategic proximity can be significantly weakened, yet represent a wide range of economic situations. The current approach combines conditions on the payoff functions with conditions on the prior which makes that the class of payoff functions for which the results hold can be fairly general. In particular, nontrivial effects of higher-order beliefs are not ruled out a priori, as illustrated in Section 2: low-probability events can have a large impact on strategic outcomes through players' higher-order beliefs.

As noted above, anonymous local games are characterized by joint conditions on payoff functions and priors. The payoff functions determine the strength of the payoff interdependencies, while the prior specifies the probability with which each of the types affect a given type's payoff. The current results critically rely on this. That the sensitivity of predictions depend on the prior (in addition to the payoff structure) is in line with the results of [18], obtained in a very different setting. In the context of coordination games played on a network, Morris [18] studies the conditions on the network structure under which an action that is played initially by a small subset of players can spread to the rest of the population under myopic best-response dynamics. Intuitively, a game on a network such that players readily switch actions if a small fraction of players does so is conceptually similar to an incomplete-information game with a prior such that equilibrium behavior is sensitive to small perturbations of beliefs (also see [17]). Morris [18] shows that interaction structures differ greatly in their sensitivity to a small change in behavior by players, just like the strength of the conditions for strategic proximity depends on the prior.

²In particular, see [4, 6, 11, 16].

³Kalai [12] also focuses on a restricted class of games to obtain more robust predictions, albeit in a different context. Weinstein and Yildiz [24] derive sufficient conditions on payoff functions and best-response correspondences so that perturbations of beliefs at high orders have negligible impact on equilibrium actions.




Figure 1: (a) The state $\omega_{m,k}$; (b) the state $\tilde{\omega}_{m,k}$. In both figures, the players are labeled 1, 2, 3, 4, 5, and their types are $t_{m-1}, t_{m-1}, t_m, t_m, t_m$, respectively; player labels are omitted for clarity. An arrow from a player i to j means that i can affect j 's payoffs, but not conversely, and a dashed line between i and j means that they can both affect the other's payoffs.

It is important to emphasize, though, that the results of [18] cannot easily be translated to the current setting. Most importantly, in the games studied in [18], the best response of a player is solely determined by the fraction of his neighbors taking an action. By contrast, the best-response correspondences considered here can be completely general. It is all the more striking that certain insights, such as the importance of cohesiveness, do carry over to this much richer setting.

The outline of this paper is as follows. Section 2 provides two examples to illustrate the different effects perturbations can have on strategic outcomes. Section 3 formally introduces the class of anonymous local games and defines strategic proximity. It also proposes a notion of closeness of priors which does not refer to any strategic aspects of the situation. Section 4 shows that the topology induced by this measure fully characterizes strategic proximity, and compares the current conditions to the conditions for strategic proximity in all Bayesian games. Section 5 illustrates how the results can be applied in the context of network games. Section 6 concludes.

2 Examples

In this section I provide an example of a game with type-dependent and heterogeneous externalities to illustrate the different effects small perturbations of prior beliefs can have. There are five players, labeled 1, 2, 3, 4 and 5, with types t_0, t_1, \dots . The payoff to a player potentially depends on the actions and types of only a subset of players. Say that a player i is a *neighbor* of another player j if i 's action or type affects j 's payoffs. Which players are neighbors depends on the state of the world. States can thus be represented by directed networks, with the (directed) edges specifying the payoff dependencies.

There are two actions, a_S and a_R . The payoff of a_S to a player is always zero. The payoff of action a_R depends on the actions and types of a players' neighbors. Specifically, the payoff of

a_R to a player i of type t_m equals:

- $b > 0$ if i has two neighbors, both have type t_m , and both play a_R ;
- $b > 0$ if i has four neighbors, two of which have type t_m , two have type t_{m+1} , and all neighbors play a_R ;
- $-c < 0$ otherwise.

Action a_S is thus a safe action, while action a_R is risky. Assume that $3b - 4c < 0$.

Depending on players' beliefs over the payoff interdependencies, small perturbations of their beliefs can have very different effects, as the following two examples demonstrate:

Example 2.1 (Containment) The state space is $\Omega = \{\omega_{m,k} : m \in \mathbb{N}, k = 1, \dots, 10\}$. In state $\omega_{m,k}$, two players have type t_{m-1} , and three players have type t_m ; the index k keeps track of the $\binom{5}{3} = 10$ different configurations of players.⁴ A state $\omega_{m,k}$ thus fixes an interaction structure; see Figure 1(a).⁵

Players have a common prior μ^* ; the prior probability of state $\omega_{m,k}$ is

$$\mu^*(\omega_{m,k}) = \left(\frac{1}{10}\right) \left(\frac{1}{2^m}\right).$$

A player of type t_m assigns conditional probability $\frac{4}{7}$ to the states in which the other players have type t_{m-1} and t_m , and conditional probability $\frac{3}{7}$ to the states in which the other players have type t_m and t_{m+1} . Consider the symmetric strategy profile σ^* in which all players choose the risky action a_R , regardless of their type. Conditional on a player's type, his expected payoffs are $\frac{3}{7}b - \frac{4}{7}c$, which is strictly positive by assumption. Since no player can gain by deviating, σ^* is an equilibrium, and the ex ante expected payoffs are $\frac{3}{7}b - \frac{4}{7}c$ for each player.

Now consider the sequence of priors μ^1, μ^2, \dots such that

$$\mu^\ell(\omega_{m,k}) = \frac{1}{M^\ell} \left(\frac{1}{10}\right) \left(\frac{1}{2^m}\right)$$

if $\ell \leq m$, and $\mu^\ell(\omega_{m,k}) = 0$ otherwise, with $1/M^\ell$ a normalization factor. Is the strategy profile in which all players choose the risky action a_R for every possible type still an equilibrium under the perturbed prior μ^ℓ ? Under the perturbed prior μ^ℓ , a player of type t_ℓ assigns probability 1 to the event that he has three neighbors of type $t_{\ell-1}$, and one of type t_ℓ . Playing a_R thus gives him an expected payoff of $-c < 0$. Hence, in any equilibrium, a player of type t_ℓ chooses a_S , and obtains a payoff of 0. I claim that the strategy profile σ^ℓ in which all players play the safe

⁴For instance, in one of the possible configurations of players, players 1 and 2 have type t_m , and the other players have type t_{m-1} . The labeling of configurations (by k) is fixed arbitrarily.

⁵In the formal treatment, more general state spaces are considered. This should not create any confusion.

action a_S if their types is t_ℓ and the risky action a_R otherwise is an equilibrium. To see this, note that the payoffs of a_R to a player of type $t_{\ell-1}$ are not affected by which action a player of type t_ℓ takes. Hence, playing a_R is still optimal for a player of that type. Given that players of type $t_{\ell-1}$ choose a_R , the expected payoffs of a_R under the prior μ^ℓ to players of type $t_0, \dots, t_{\ell-2}$ are the same as under the prior μ^* , so that the risky action a_R is optimal for players with these types. As ℓ goes to infinity, ex ante expected payoffs under this equilibrium converge to $\frac{3}{7}b - \frac{4}{7}c$, the expected payoff of the strategy profile σ^* in which all players choose the risky action for every possible type. Consequently, even if σ^* is no longer an equilibrium under the perturbed prior μ^ℓ , there is a symmetric equilibrium σ^ℓ that gives approximately the same payoffs as σ^* when ℓ goes to infinity. \triangleleft

While the effect of perturbing some players beliefs was contained in the previous example, the next example shows that a similarly small change can have a large effect on strategic outcomes in a slightly different setting:

Example 2.2 (Contagion) Consider the state space $\tilde{\Omega} = \{\tilde{\omega}_{m,k} : m \in \mathbb{N}, k = 1, \dots, 10\}$. In $\tilde{\omega}_{m,k}$, there are again two players of type t_{m-1} , and three players of type t_m , with k labeling the different assignments of players to types, as before. The difference between $\tilde{\omega}_{m,k}$ and $\omega_{m,k}$ is which players influence which players' payoffs. As in the previous example, players of the same type affect each others' payoffs, but now it is the players of type t_{m-1} whose payoffs depend on the actions of the players of the other type, rather than the other way around; see Figure 1(b). Define the prior $\tilde{\mu}^*$ by

$$\tilde{\mu}^*(\tilde{\omega}_{m,k}) = \left(\frac{1}{10}\right) \left(\frac{1}{2^m}\right).$$

As in Example 2.1, there is an equilibrium $\tilde{\sigma}^*$ in which all players choose the risky action a_R , for any type they may have, so that their ex ante expected payoffs are $\frac{3}{7}b - \frac{4}{7}c > 0$, as before. Now take a sequence of perturbed priors $\tilde{\mu}^1, \tilde{\mu}^2, \dots$, defined by

$$\tilde{\mu}^\ell(\tilde{\omega}_{m,k}) = \frac{1}{M^\ell} \left(\frac{1}{10}\right) \left(\frac{1}{2^m}\right)$$

if $m \leq \ell$, and $\tilde{\mu}^\ell(\tilde{\omega}_{m,k}) = 0$ otherwise, with M^ℓ a normalization factor. While in Example 2.1 there was an equilibrium “close” to the original equilibrium σ^* when the perturbations became arbitrarily small (in terms of prior probabilities), this is not the case here.⁶ Indeed, choosing the risky action a_R gives a player of type t_ℓ a payoff of $-c < 0$ (independent of other players' actions), while the safe action guarantees him a payoff of 0. Hence, in any equilibrium, players of type t_ℓ play a_S . But that means that choosing a_R gives a player of type $t_{\ell-1}$ a payoff of $-c$, so that players of that type will play a_S in equilibrium, leading players of type $t_{\ell-2}$ to choose a_S ,

⁶In the formal treatment, I allow for approximate equilibria. This would not affect the argument here.

and so on. Hence, in equilibrium, all players choose the safe action a_S for every possible type, giving ex ante expected payoffs of 0. There is thus no equilibrium under the prior $\tilde{\mu}^\ell$ such that players receive payoffs that are similar to those under $\tilde{\sigma}^*$, even for ℓ large. \triangleleft

In both cases, the perturbed priors are similar to the original prior in terms of the ex ante probabilities they assign to events. However, even when the perturbations in terms of prior probabilities become arbitrarily small, some types assign very different conditional probabilities to certain events under the perturbed priors. In the first case, the effect of a player with perturbed conditional beliefs choosing a different equilibrium action is “contained” because this player has a limited impact on other players’ payoffs. In the second example, however, the actions of players with perturbed conditional beliefs do have a substantial effect on other players’ payoffs, leading these players to choose different actions in equilibrium, and so on. Section 4 investigates the conditions under which the effects of perturbations on equilibrium outcomes can be contained, after the necessary concepts have been introduced in the next section.

3 Framework

3.1 Anonymous local games

An anonymous local game is a game in which a players’ payoff may depend on only a subset of players. Also, players are assumed to be symmetric ex ante, and payoff functions are anonymous. These properties make that the equilibria of the game can be analyzed without making reference to the number of players, so that games with an unbounded number of players can be analyzed. To allow for this, it is convenient to work with a countable set of potential players. Formally, there is a set of (*candidate*) *players* $N_c := \mathbb{Z}$ who face uncertainty about payoffs. Payoff uncertainty is represented by a state space $\Omega = \mathcal{G} \times \mathcal{S} \times \mathcal{C}$, where \mathcal{G} is the set of interaction structures, \mathcal{C} is the collection of sets of *active players* or *agents*, and the space \mathcal{S} represents other payoff-relevant uncertainty. The sets \mathcal{G} and \mathcal{C} are discussed in more detail below.

Nature draws a pair (g, s) from $\mathcal{G} \times \mathcal{S}$ according to a *common prior* μ . An interaction structure g can be represented by a directed network, so that the set \mathcal{G} of interaction structures can be taken to be the (countable) set of all directed networks with a finite number of vertices.⁷ Conditional on the event that the number of vertices of the interaction structure g equals n , a set N of n agents is drawn according to an improper uniform prior from the collection $\mathcal{C}_n \subset \mathcal{C}$ of subsets of n candidate players, where \mathcal{C}_n is the collection of subsets $\{m, m + 1, \dots, m + n - 1\}$ of n candidate players with consecutive labels, with $m \in N_c = \mathbb{Z}$. Active players are associated

⁷There cannot be an edge from an agent to himself. Also, I do not allow for multiple edges between a given pair of agents, or for weighted edges.

with a vertex of the interaction structure g in such a way that each active player is equally likely to occupy a certain position in the interaction structure.⁸ A (directed) edge from an agent i to another agent j in the directed network representing g means that j 's payoff potentially depends on the action or type of agent i (depending on the payoff function). In that case, i is said to be a *neighbor* of j . The set of neighbors of j in state $\omega = (g, s, N)$ is denoted by $N_j(g, N)$.

The common prior μ on $\mathcal{G} \times \mathcal{S}$ and the improper uniform prior on the collection \mathcal{C} of sets of active players give a prior η on the state space $\Omega = \mathcal{G} \times \mathcal{S} \times \mathcal{C}$. Even though the prior η is improper, the measure η_N on Ω obtained from η by conditioning on the event that the set of active players is N is a well-defined probability measure; see e.g. [5, pp. 190–191].⁹ The expected number of active players is assumed to be finite.

The payoffs to each active player depend on his own action and type and the actions and types of his neighbors. Each candidate player is endowed with a finite set of actions A , the same for each player. The type of player i in state $\omega = (g, s, N)$ is denoted by $\tau_i(g, s, N)$ or sometimes $\tau_i(\omega)$ for brevity. A player's type specifies whether he is selected to play, but otherwise contains information only on the interaction structure and the payoff parameter. The set of types T is the same for each candidate player and is assumed to be countable. The set of types that know that they participate is denoted T_a , the set of *active types*.

The payoffs to an active player $i \in N$ in state $\omega = (g, s, N)$ given the action profile $a = (a_j)_{j \in N}$ are:

$$u_i(a, g, s, N) = v_{|N_i(g, N)|}(a_i, (a_j)_{j \in N_i(g, N)}, \tau_i(g, s, N), (\tau_j(g, s, N))_{j \in N_i(g, N)})$$

if the set of neighbors $N_i(g, N)$ of i is nonempty. That is, payoffs do not depend on the identity of a player, and are a function only of the player's action and type and the actions and types of his neighbors. If player i has no neighbors, or if i is not active in state $\omega = (g, s, N)$, then $u_i(a, g, s, N) = v_0(a_i, t_i)$, i.e., the payoffs to i are a function only of his own action and type.

The payoff functions v_k , $k = 1, 2, \dots$, are assumed to be *anonymous*. That is, for each $b \in A$, $(a_1, \dots, a_k) \in A^k$, $t \in T$, $(t_1, \dots, t_k) \in T_a^k$,

$$v_k(b, (a_1, \dots, a_k), t, (t_1, \dots, t_k)) = v_k(b, (a_{\pi(1)}, \dots, a_{\pi(k)}), t, (t_{\pi(1)}, \dots, t_{\pi(k)})) \quad (1)$$

for any permutation π of $1, \dots, k$. While the anonymity condition (1) rules out that an agent's payoffs depend on the identity of a neighbor who takes a certain action or has a certain type, it allows for a rich class of payoff interdependencies. For instance, an agent's payoff may differ depending on whether a neighbor of type t or type t' takes a given action a . The functions v_0, v_1, \dots are referred to as *(local) payoff functions*. The profile $v := (v_0, v_1, \dots)$ of local payoff

⁸The procedure that associates active players to vertices is assumed to be fixed across states.

⁹Also see [14] for a discussion of the implications of allowing for an unbounded random number of players.

functions is bounded by B if

$$|v_k(a^{(k+1)}, t, t^{(k)}) - v_k(b^{(k+1)}, t, t^{(k)})| \leq B \quad \text{and} \quad |v_k(a^{(k+1)}, t, t^{(k)})| \leq B$$

for any k , action profiles $a^{(k+1)}, b^{(k+1)} \in A^{k+1}$, type $t \in T_a$, and type profile $t^{(k)} \in T_a^k$. A profile of payoff functions is *bounded* if it is bounded by some $B \in \mathbb{R}$.

I also assume that the types of active players are *exchangeable*. That is, for any set $N \subset N_c$ of agents, any set of distinct players i_1, \dots, i_m in N , any active types $t_1, \dots, t_m \in T_a$, and any permutation π of $1, \dots, m$,

$$\eta_N(\{(g, s, M) \in \Omega : \tau_{i_1}(g, s, M) = t_1, \dots, \tau_{i_m}(g, s, M) = t_m\}) = \eta_N(\{(g, s, M) \in \Omega : \tau_{i_{\pi(1)}}(g, s, M) = t_1, \dots, \tau_{i_{\pi(m)}}(g, s, M) = t_m\}),$$

where η_N is the probability measure on the state space conditional on the event that the set of active players is N .¹⁰ Together, the exchangeability and anonymity assumptions make that players are *symmetric* from an ex ante perspective.

It will be convenient to define the type profile of an agent's neighbors without making reference to their identity. Label the types in T_a as t^1, t^2, \dots . If in state $\omega = (g, s, N)$ agent $i \in N$ has a nonempty set of neighbors $N_i(g, N) := \{j_1, \dots, j_k\}$, and the types of these neighbors are given by $t_{j_1}, \dots, t_{j_k} \in T_a$, then the *neighbor type profile* $\theta_i = \theta_i(g, s, N)$ of i in (g, s, N) is simply the tuple $(t_{\rho(j_1)}, \dots, t_{\rho(j_k)})$ where ρ is a permutation of j_1, \dots, j_k that places the types in a non-decreasing order (according to the labeling of the elements of T_a). Otherwise, if agent i has no neighbors, set $\theta_i := (0)$. For future reference, let $\Theta^{(k)} \subset T_a^k$ be the set of neighbor type profiles of dimension k , and let $\Theta := \bigcup_k \Theta^{(k)}$ be the set of all neighbor type profiles.

Since an agent's type does not contain any information on the (random) set of active players other than that he is selected to play, it follows that the probability that an arbitrary subset of active players has a given type or neighbor type profile is independent of the identity of the active players. Because the interaction structure completely determines the number of agents, it suffices to consider the (proper) prior μ on the set of interaction structures when referring to agents' beliefs about types and (neighbor) type profiles. Formally, the probability $\mu(t)$ that an arbitrary agent has type t is:

$$\mu(t) := \frac{\sum_{(g,s) \in \mathcal{G} \times \mathcal{S}} n_g \mu(\{(g, s) : \tau_1(g, s) = t\})}{\sum_{(g',s') \in \mathcal{G} \times \mathcal{S}} n_{g'} \mu(\{(g', s')\})},$$

where, with some abuse of notation, $\tau_1(g, s)$ is the type of the agent associated with vertex 1 in g (which clearly does not depend on the set of active players), n_g is the number of vertices

¹⁰This is with some abuse of notation, as the probability that a given active player is associated with a given vertex is not given by η . Because the protocol that matches agents with vertices is fixed across states, no confusion should result.

(players) in interaction structure g and where the weight $n_g / (\sum_{(g',s)} n_{g'} \mu(\{(g', s')\}))$ (the number of players in g relative to the average number of players) is applied to account for the fact that conditional on being selected to play, an agent assigns a higher probability to participating in a game with many players [20].¹¹ Finally, for notational convenience assume that, with probability one, each agent has at least one neighbor. Denote the class of priors μ on $\mathcal{G} \times \mathcal{S}$ that satisfy the above conditions by \mathcal{M} . Throughout this paper, an anonymous local game with prior $\mu \in \mathcal{M}$ on $\mathcal{G} \times \mathcal{S}$ and local payoff function profile v is denoted by (μ, v) .

3.2 Strategies and equilibrium

As is standard, a (mixed) *strategy* of a player i is a function $\sigma_i : T \rightarrow \Delta(A)$, with $\sigma_i(t_i)(a_i)$ the probability that i chooses action a_i when his type is t_i . A *strategy profile* $\sigma = (\sigma_i)_{i \in N_c}$ specifies a strategy for each candidate player i . A strategy profile is *symmetric* if $\sigma_i = \sigma_j$ for all $i, j \in N_c$. As always, σ_{-i} denotes the strategy profile $(\sigma_j)_{j \in N_c \setminus \{i\}}$ of players other than i .

Fix an active type $t_i \in T_a$ such that t_i has positive probability under the common prior μ . The *interim expected payoff* of action $a_i \in A$ to an active player i of type t_i is given by

$$\varphi_i(a_i, \sigma_{-i}; t_i, \mu) := \sum_{\substack{\omega=(g,s,N) \in \Omega, \\ D \subseteq N \setminus \{i\}}} \eta_i(\omega, D \mid t_i) v_{|D|}(a_i, (\sigma_j(\tau_j(\omega)))_{j \in D}, \tau_i(\omega), (\tau_j(\omega))_{j \in D}),$$

where $\eta_i(\omega, D \mid t_i)$ is the conditional probability (given that i is active and has type $t_i \in T_a$) that the state is $\omega = (g, s, N)$ and that i 's set of neighbors is D . If σ is a symmetric strategy profile, this can be written as:

$$\varphi_i(a_i, \sigma_{-i}; t_i, \mu) = \sum_{\substack{k \in \mathbb{N}, \\ \theta = (\theta_1, \dots, \theta_k) \in \Theta^{(k)}}} \mu(\theta \mid t_i) v_k(a_i, (\sigma_j(\theta_j))_{j=1, \dots, k}, t_i, \theta),$$

where $\mu(\theta \mid t_i)$ is the conditional probability that an agent has neighbor type profile $\theta = (\theta_1, \dots, \theta_k)$ conditional on his type being $t_i \in T_a$. For a candidate player i who is not selected into the game, or for a type t_i that has zero probability under μ , simply set $\varphi_i(a_i, \sigma_{-i}; t_i, \mu) = 0$ for any a_i and σ_{-i} .

For $\varepsilon \geq 0$, a strategy profile σ is an ε -*equilibrium* of an anonymous local game (μ, v) if for each candidate player $i \in N_c$ with an active type $t_i \in T_a$ such that $\mu(t_i) > 0$, it holds that

$$\varphi_i(a_i, \sigma_{-i}; t_i, \mu) \geq \varphi_i(b_i, \sigma_{-i}; t_i, \mu) - \varepsilon$$

¹¹This is easy to see. Suppose that the number of active players is 2 or 2000 with equal probability. *Conditional on being selected to play*, an agent assigns probability $(\frac{1}{2} \cdot 2000) / (\frac{1}{2} \cdot 2000 + \frac{1}{2} \cdot 2) > \frac{1}{2}$ to participating in the game with 2000 agents; see [20] for an extensive discussion.

for each $a_i \in A$ with $\sigma_i(t_i)(a_i) > 0$ and for any $b_i \in A$. That is, in an ε -equilibrium, an active player can gain at most ε from deviating; a 0-equilibrium is referred to as an *equilibrium*. An ε -equilibrium is *symmetric* if the strategy profile is symmetric.

Proposition 3.1 *If an anonymous local game (μ, v) is bounded, it has a symmetric equilibrium.*

The proof can be found in the online appendix. The proof relies on the assumption that the expected number of agents is finite. Together with the assumption that the profile of local payoff functions is bounded, this ensures that the payoff functions are continuous.¹²

It will also be useful to define a player’s expected payoffs conditional on being selected into the game, but before he learns his type. The *expected payoffs* of a player i are obtained from the interim expected payoffs in the usual way, by averaging over all possible active types $t_i \in T_a$, weighted by the probability $\mu(t_i)$ of that type:

$$\Phi_i(\sigma; \mu) = \sum_{t_i \in T_a} \mu(t_i) \sum_{a_i \in A} \sigma_i(t_i)(a_i) \varphi_i(a_i, \sigma_{-i}; t_i, \mu).$$

Hence, an anonymous local game is essentially a Bayesian game with a countably infinite player set which satisfies a number of symmetry and anonymity assumptions. In particular, if the common prior assigns probability one to an interaction structure with a fixed number of agents n , and the set of active players N becomes common knowledge among the agents, then an anonymous local game can be seen as a standard Bayesian game with a finite player set. To avoid tedious repetition, the qualification “conditional on being selected into the game” when referring to an agent’s beliefs will be omitted in the following.

3.3 Strategic closeness

The aim of the paper is to define a measure on priors such that if two priors in \mathcal{M} are close according to this measure, then players can expect similar outcomes in equilibrium under the two priors in any anonymous local game with bounded payoffs. In that case, the priors are strategically close. Here I make precise what is meant by strategic proximity; Section 3.4 proposes a measure on priors that does not make references to the strategic features of the situation. This measure will be used in Section 4 to characterize strategic proximity in anonymous local games.

Consider two anonymous local games (μ, v) and (μ', v) such that v is a bounded profile of local payoff functions. What does it mean for an agent to obtain similar equilibrium payoffs in the games (μ, v) and (μ', v) ? Adapting a notion proposed by [11] for Bayesian games to the context

¹²The proof builds on an insight of [21], who showed that payoff functions in a complete-information game with an arbitrary player set are continuous if they can be approximated by payoff functions that only depend on the actions of a finite subset of players. If types are independent, the result of [1] for Bayesian games with a countably infinite player set can be used to obtain existence of a (possibly asymmetric) Bayesian equilibrium.

of anonymous local games, I say that μ and μ' are strategically close if for any bounded profile v , for any symmetric equilibrium σ of the game (μ, v) , there exists a symmetric approximate equilibrium σ' of the game (μ', v) such that expected payoffs under σ and σ' are close, and vice versa.

Formally, let $\mu, \mu' \in \mathcal{M}$ be common priors, and let $v := (v_k)_{k \in \mathbb{N}}$ be a profile of local payoff functions. Note that if all opponents play according to a symmetric strategy profile σ , the interim expected payoffs and expected payoffs of a player do not depend on his identity. That is, for any $i, j \in N_c$ and $t \in T_a$, it holds that $\varphi_i(\cdot, \sigma_{-i}; t, \mu) = \varphi_j(\cdot, \sigma_{-j}; t, \mu)$ and $\Phi_i(\sigma; \mu) = \Phi_j(\sigma; \mu)$ for any strategy profile σ and prior μ , so that it is sufficient to consider a fixed player $i \in N_c$. For each $\varepsilon \geq 0$, define

$$\chi(\mu, \mu'; v, \varepsilon) := \sup_{\sigma \in \mathcal{E}^0(\mu, v)} \inf_{\sigma' \in \mathcal{E}^\varepsilon(\mu', v)} |\Phi_i(\sigma; \mu) - \Phi_i(\sigma'; \mu')|,$$

where $\mathcal{E}^\varepsilon(\mu, v)$ is the set of symmetric ε -equilibria of (μ, v) , so that $\mathcal{E}^0(\mu, v)$ denotes the set of symmetric equilibria of the game. To obtain a measure that is symmetric in μ and μ' , let

$$\chi^*(\mu, \mu'; v, \varepsilon) := \max \{ \chi(\mu, \mu'; v, \varepsilon), \chi(\mu', \mu; v, \varepsilon) \}.$$

If $\chi^*(\mu, \mu'; v, \varepsilon)$ is small for two priors $\mu, \mu' \in \mathcal{M}$ for any $\varepsilon > 0$ and any bounded profile of local payoff functions v , then for *any* symmetric equilibrium of the anonymous local game (μ, v) , there exists *some* approximate symmetric equilibrium of the game (μ', v) such that expected payoffs are similar, for any $\varepsilon > 0$ and any bounded profile of payoff functions v ; conversely, for any symmetric equilibrium in (μ', v) , there is an approximate symmetric equilibrium in (μ, v) that gives similar expected payoffs. If that is the case, μ and μ' are said to be *strategically close* on the class of anonymous local games. If $\chi^*(\mu, \mu^\ell; v, \varepsilon) \rightarrow 0$ for a sequence of priors $(\mu^\ell)_{\ell \in \mathbb{N}}$ for any $\varepsilon > 0$ and any bounded profile v , the sequence $(\mu^\ell)_{\ell \in \mathbb{N}}$ *converges strategically* to the prior μ ; this is the case for instance for the sequence of priors $(\mu^\ell)_{\ell \in \mathbb{N}}$ and the prior μ^* in Example 2.1.

3.4 A topology on priors

The purpose of this section is to define a notion of “closeness” of priors that does not make reference to strategic features, such as payoff functions. In the next section, I show that closeness according to this measure is a necessary and sufficient condition for strategic proximity in anonymous local games.

To characterize the closeness of priors, I formulate two conditions. The first captures weak convergence of priors, which is equivalent to uniform weak convergence in the current context (recall that the state space is countable). Let $\mu(F)$ be the probability that the type and neighbor type profile of an agent lie in $F \subseteq T_a \times \Theta$, and define

$$d_0(\mu, \mu') := \sup_{F \subseteq T_a \times \Theta} |\mu(F) - \mu'(F)|. \quad (2)$$

If μ and μ' are close according to this measure, then they assign similar prior probabilities to all relevant events. The perturbed priors in Examples 2.1 and 2.2 converge to the priors μ^* and $\tilde{\mu}^*$, respectively, in this sense.

The second condition concerns the conditional beliefs of agents. To state this condition, I define a belief operator that is suitable for the case of anonymous local games.¹³ The *local p -belief operator* B_μ^p associates with each set of types the subset of types such that with conditional probability at least p , their payoffs depend only on the actions and types of agents with types in that set. Formally, for any subset $U \subseteq T_a$ of active types, let $\Xi(U) := \bigcup_k (U^k \cap \Theta^{(k)})$ be the set of neighbor type profiles such that each neighbor's type belongs to the set of types U . For any $t \in T_a$ such that $\mu(t) > 0$, let $\mu(\Xi(U) \mid t)$ be the probability that the types of the neighbors of an agent belong to U , given that his type is t . Then,

$$B_\mu^p(U) := \{t \in U : \mu(t) > 0 \Rightarrow \mu(\Xi(U) \mid t) \geq p\}.$$

Note that $B_\mu^p(U)$ includes the types in U that have probability zero. The online appendix discusses a number of properties of the local belief operator.

By definition, $B_\mu^p(U) \subseteq U$. If in addition

$$B_\mu^p(U) \supseteq U,$$

then the set of types U is said to be *p -cohesive* (under μ). This term is motivated by the following: If a set of types U is p -cohesive, then with high conditional probability the payoffs to each player with a type in U depend only on players with types in that set such that with high conditional probability, the payoffs of those players depend only on players with types in that set, and so on.

Define $[B_\mu^p]^1(U) := B_\mu^p(U)$ and, for each $\ell \in \mathbb{N}$, let $[B_\mu^p]^{\ell+1} = B_\mu^p \circ [B_\mu^p]^\ell$. Let

$$C_\mu^p(U) := \bigcap_{\ell \in \mathbb{N}} [B_\mu^p]^\ell(U).$$

Intuitively, $C_\mu^p(U)$ is the set of types in U such that with conditional probability at least p , their payoffs depend only on types in U such that with conditional probability at least p , their payoffs . . . depend only on types in U , for any number of iterations. Lemma A.1 in the appendix shows that $C_\mu^p(U)$ is indeed p -cohesive.

To use the local p -belief operator to formulate conditions on agents' conditional beliefs, denote the conditional probability that the neighbor type profile of an agent of type t lies in $F \subseteq \Theta$ by $\mu(F \mid t)$, and define

$$T_{\mu, \mu'}^\delta := \left\{ t \in T_a : \mu(t), \mu'(t) > 0 \text{ and } \sup_{F \subseteq \Theta} |\mu(F \mid t) - \mu'(F \mid t)| \leq \delta \right\}$$

¹³ [17] and [19] propose similar operators in a different context.

to be the set of active types such that agents' conditional beliefs about their neighbors' types are within δ . Clearly, the optimal action for an agent with a type t that induces different conditional beliefs under μ and μ' (i.e., $t \notin T_{\mu, \mu'}^\delta$) may be different under the two priors. It turns out, however, that even if with high (prior) probability, an agent has a type such that his conditional beliefs about his neighbors' types are similar under μ and μ' (i.e., $t \in T_{\mu, \mu'}^\delta$), outcomes may differ considerably under the two priors, so that a stronger condition is required. Define

$$d_1(\mu, \mu') := \inf \{ \delta : \mu(\{(g, s) : \tau_1(g, s) \in C_\mu^{1-\delta}(T_{\mu, \mu'}^\delta)\}) \geq 1 - \delta \}, \quad (3)$$

where $\tau_1(g, s)$ is the type of the agent associated with vertex 1 in the interaction structure g . If $d_1(\mu, \mu')$ is small, then, with high prior probability (under μ), an agent has a type such that his conditional beliefs are similar under μ and μ' , and with high conditional probability, his payoffs depend only on players whose conditional beliefs are close under the priors, and whose payoffs depend (with high conditional probability) only on players with close conditional beliefs, and so on.¹⁴ That is, $d_1(\mu, \mu^\ell) \rightarrow 0$ for some sequence of priors $(\mu^\ell)_{\ell \in \mathbb{N}}$ whenever the set of types in T_a with arbitrarily close conditional beliefs is p -cohesive for p arbitrarily close to one, and the prior probability of this set goes to one. This condition is strictly weaker than the requirement that conditional beliefs converge uniformly.

Combining (2) and (3) gives:

$$d^*(\mu, \mu') := \max \{ d_0(\mu, \mu'), d_1(\mu, \mu'), d_1(\mu', \mu) \}.$$

It is immediate that d^* is nonnegative and symmetric. Moreover, $d^*(\mu, \mu') = 0$ if and only if $\mu = \mu'$. Because d^* need not satisfy the triangle inequality, it is not a metric. However, the measure d^* introduces a topology on the set of priors \mathcal{M} by specifying a notion of convergence, where a sequence $(\mu^\ell)_{\ell \in \mathbb{N}}$ converges to μ if and only if for any $\varepsilon > 0$, there exists $\Lambda \in \mathbb{N}$ such that $d^*(\mu^\ell, \mu) \leq \varepsilon$ for all $\ell > \Lambda$. This topology will be used in the next section to characterize strategic proximity in anonymous local games.

4 Results

4.1 Strategic proximity in anonymous local games

The previous section developed a measure d^* on the class of priors \mathcal{M} which defines a topology on \mathcal{M} which is stronger than the topology of weak convergence (of prior beliefs), but whose notion

¹⁴Note that (3) only refers to the event that some fixed agent has a type in $C_\mu^{1-\delta}(T_{\mu, \mu'}^\delta)$, not to the event that the neighbors of this agent have their type in $C_\mu^{1-\delta}(T_{\mu, \mu'}^\delta)$, or to the event that *all* agents have their type in $C_\mu^{1-\delta}(T_{\mu, \mu'}^\delta)$ [11]. The reason is that $C_\mu^{1-\delta}(U)$ is $(1-\delta)$ -cohesive: If a fixed agent has a type in $C_\mu^{1-\delta}(U)$ for some subset of types U , then with high probability all his neighbors have a type in U ; see Lemma A.2.

of convergence does not require uniform convergence of conditional beliefs. Theorem 4.1 shows that convergence of priors in this topology is a necessary and sufficient condition for strategic convergence in anonymous local games:

Theorem 4.1 *Let $\mu \in \mathcal{M}$ and let $(\mu^\ell)_{\ell \in \mathbb{N}}$ be a sequence in \mathcal{M} . Then, $\chi^*(\mu, \mu^\ell; v, \varepsilon) \rightarrow 0$ for any $\varepsilon > 0$ and any bounded profile v of local payoff functions if and only if $d^*(\mu, \mu^\ell) \rightarrow 0$.*

The proof of Theorem 4.1 follows from Lemmas 4.2–4.4. First, Lemma 4.2 shows that the conditions in the theorem are sufficient:

Lemma 4.2 *Let $\mu, \mu' \in \mathcal{M}$, and let v be a profile of payoff functions. Suppose that $d^*(\mu, \mu') \leq \delta$. Then, if σ is a symmetric equilibrium of the game (μ, v) and v is bounded by B , there exists a symmetric $5\delta B$ -equilibrium σ' of the game (μ', v) such that*

$$|\Phi_i(\sigma; \mu) - \Phi_i(\sigma'; \mu')| \leq (4 - \delta)\delta B$$

for all $i \in N_c$.

The proof can be found in Appendix A. Second, Lemma 4.3 establishes that closeness in terms of prior probabilities is a necessary condition for strategic proximity:

Lemma 4.3 *Let $\mu, \mu' \in \mathcal{M}$. If $d_0(\mu, \mu') > \delta$, then there exists a profile v of payoff functions with bound $B = 1$ and a symmetric equilibrium σ of the game (μ, v) such that for any symmetric δ -equilibrium σ' of (μ', v) ,*

$$|\Phi_i(\sigma; \mu) - \Phi_i(\sigma'; \mu')| > \delta$$

for all $i \in N_c$.

For a proof, see Appendix A. The last building block of the proof of Theorem 4.1 is Lemma 4.4. Its proof demonstrates how small perturbations of higher-order beliefs can lead to large changes in strategic outcomes.

Lemma 4.4 *Let $\mu, \mu' \in \mathcal{M}$. If $d_1(\mu, \mu') > \delta$, then there exists a profile v of payoff functions with bound $B = 3$ and a symmetric equilibrium σ of the game (μ, v) such that for any symmetric δ -equilibrium σ' of the game (μ', v) ,*

$$|\Phi_i(\sigma; \mu) - \Phi_i(\sigma'; \mu')| > \delta^2,$$

for all $i \in N_c$.

Proof. By assumption, $d_1(\mu, \mu') > \delta$, so that

$$\mu(\{(g, s) \in \mathcal{G} \times \mathcal{S} : \tau_1(g, s) \in C_\mu^{1-\delta}(T_{\mu, \mu'}^\delta)\}) > 1 - \delta$$

or

$$\mu'(\{(g, s) \in \mathcal{G} \times \mathcal{S} : \tau_1(g, s) \in C_{\mu'}^{1-\delta}(T_{\mu, \mu'}^\delta)\}) > 1 - \delta. \quad (4)$$

Without loss of generality, assume that (4) holds. By definition, for every $t \in T_a \setminus T_{\mu, \mu'}^\delta$, there exists a set of neighbor type profiles $F_t^* \subseteq \Theta$ such that

$$\mu'(F_t^* | t) - \mu(F_t^* | t) > \delta,$$

where $\mu'(F_t^* | t)$ and $\mu(F_t^* | t)$ are the conditional probabilities under μ' and μ , respectively, that an agent's neighbor type profile belongs to F_t^* , given that his type is t .

Label the actions as $A = \{b^1, b^2, \dots, b^m\}$, and let payoffs be defined as follows.¹⁵ For each $t \in T_a$, $k \in \mathbb{N}$, $a = (a_1, \dots, a_k) \in A^k$ and $\theta \in \Theta^{(k)}$, let

$$v_k(b^1, a, t, \theta) = 0, \\ v_k(b^2, a, t, \theta) = \begin{cases} 2 & \text{if } t \in T_{\mu, \mu'}^\delta \text{ and } a_j = b^2 \text{ for some } j = 1, \dots, k; \\ -\delta & \text{if } t \in T_{\mu, \mu'}^\delta \text{ and } a_j = b^1 \text{ for all } j = 1, \dots, k; \\ 1 - \mu(F_t^* | t) & \text{if } t \in T_a \setminus T_{\mu, \mu'}^\delta \text{ and } \theta \in F_t^*; \\ -\mu(F_t^* | t) & \text{if } t \in T_a \setminus T_{\mu, \mu'}^\delta \text{ and } \theta \notin F_t^*; \end{cases}$$

and for $\ell = 3, \dots, m$, set

$$v_k(b^\ell, a, t, \theta) = -2.$$

That is, action b^1 always gives a payoff of 0, regardless of the actions and types of an agent and his neighbors. For agents of type $t \in T_{\mu, \mu'}^\delta$, action b^2 is only profitable if there is at least one neighbor who also takes action b^2 . By contrast, the payoff of b^2 to an agent of type $t \in T_a \setminus T_{\mu, \mu'}^\delta$ only depends on his neighbor type profile θ . All other actions than b^1 and b^2 are strictly dominated.

The game (μ, v) has a symmetric equilibrium σ in which all agents choose action b^1 for any type in T_a , so that the expected payoffs are $\Phi_i(\sigma; \mu) = 0$ for all $i \in N_c$. Now consider the game (μ', v) . By definition, for each type $t \in T_a \setminus T_{\mu, \mu'}^\delta$, $\mu'(F_t^* | t) - \mu(F_t^* | t) > \delta$. The interim expected payoffs of playing b^2 for an active player i of type $t \in T_a \setminus T_{\mu, \mu'}^\delta$ are then

$$\varphi_i(b^2, \sigma_{-i}; t, \mu') = \mu'(F_t^* | t)(1 - \mu(F_t^* | t)) - (1 - \mu'(F_t^* | t))\mu(F_t^* | t) > \delta$$

for every strategy profile σ_{-i} . It follows that in any δ -equilibrium, agents of type $t \in T_a \setminus T_{\mu, \mu'}^\delta$ choose action b^2 . Next consider a type $t \in T_{\mu, \mu'}^\delta$. If the conditional probability that the types of all neighbors of an agent of type t are in $T_{\mu, \mu'}^\delta$ is smaller than $1 - \delta$, i.e., if $\mu(\Xi(T_{\mu, \mu'}^\delta) | t) < 1 - \delta$, then, with conditional probability at least δ , an agent of type t has at least one neighbor who plays b^2 . Hence, the interim expected payoffs of b^2 to an agent that has such a type are at least

$$\delta \cdot 2 - (1 - \delta) \cdot \delta > \delta,$$

¹⁵This game is based on the ‘infection game’ of [11].

so that in any δ -equilibrium, agents of type $t \in T_{\mu, \mu'}^\delta$ such that $\mu(\Xi(T_{\mu, \mu'}^\delta) | t) < 1 - \delta$ will play b^2 . By a similar argument, agents of type $t \in T_{\mu, \mu'}^\delta$ such that $\mu(\Xi(B_{\mu'}^{1-\delta}(T_{\mu, \mu'}^\delta)) | t) < 1 - \delta$ will play b^2 in any δ -equilibrium, and so on. Consequently, all agents of type $t \in T_{\mu, \mu'}^\delta$ such that $\mu(\Xi(C_{\mu'}^{1-\delta}(T_{\mu, \mu'}^\delta)) | t) < 1 - \delta$ will play b^2 in any δ -equilibrium.

By (4), the probability that an agent has a type $t \in T_a \setminus C_{\mu'}(T_{\mu, \mu'}^\delta)$ is greater than δ . By Lemma A.1 in the appendix, the set $C_{\mu'}^{1-\delta}(T_{\mu, \mu'}^\delta)$ is $(1 - \delta)$ -cohesive, so that the probability that an agent has a type $t \in T_{\mu, \mu'}^\delta$ such that $\mu(C_{\mu'}^{1-\delta}(T_{\mu, \mu'}^\delta) | t) < 1 - \delta$ is greater than δ . This implies that in any δ -equilibrium σ' of (μ', v) , players' expected payoffs are greater than δ^2 , and the result follows. \square

The proof of Theorem 4.1 follows directly from Lemmas 4.2–4.4:

Proof. (If) Let v be a profile of payoff functions. By Lemma 4.2, if v is bounded by B , and for $\ell \in \mathbb{N}$ such that $5Bd^*(\mu, \mu^\ell) \leq \varepsilon$,

$$\chi^*(\mu, \mu^\ell; v, \varepsilon) \leq (4 - d^*(\mu, \mu^\ell))d^*(\mu, \mu^\ell)B.$$

Hence, for all profiles of payoff functions v that are bounded and for all $\varepsilon > 0$, if $d^*(\mu, \mu^\ell) \rightarrow 0$, then $\chi^*(\mu, \mu^\ell; v, \varepsilon) \rightarrow 0$.

(Only if) Let $\mu, \mu' \in \mathcal{M}$ and fix $\delta \in [0, 1)$. If $d_0(\mu, \mu') > \delta$ or $d_1(\mu, \mu') > \delta$, then, by Lemmas 4.3 and 4.4, there is a profile of payoff functions v bounded by $B = 3$ and a symmetric equilibrium σ of (μ, v) such that for any symmetric δ -equilibrium σ' of (μ', v) , $|\Phi_i(\sigma; \mu) - \Phi_i(\sigma'; \mu')| > \delta^2$. \square

To gain a better insight into the conditions for strategic proximity in anonymous local games, refer back to the examples in Section 2. As noted above, the perturbed priors converge to the original priors in terms of the prior probabilities assigned to events. Also, in both cases there is a single type whose conditional beliefs under the perturbed prior are very different than its belief under the original prior, and this type has vanishingly small probability in the limit. The key difference between the two examples is that in the first case, a switch to the safe action by a player of type t_ℓ had no effect on the optimal action of other players, while in the second case, such a change led all types to choose the safe action. That is, even if a player believes that all his neighbors have conditional beliefs close to their conditional beliefs under the original prior, believe that his neighbors believe this, and so on, up to high order, it is optimal for him to choose the safe action if he believes that some his neighbors believe that...some of their neighbors believe that the conditional beliefs of their neighbors are very different from their beliefs under the original prior.

4.2 Comparison to general Bayesian games

How do the conditions for strategic proximity in anonymous local games in Theorem 4.1 compare to the corresponding conditions for general Bayesian games? As noted in Section 3, if the number of players is fixed and the set of (active) players is common knowledge, then an anonymous local game is essentially a Bayesian game that satisfies a number of anonymity and symmetry properties. Since the anonymous local games with a commonly known player set form a subclass of the Bayesian games, the conditions derived in the previous section cannot be stronger than the corresponding conditions for strategic proximity for all Bayesian games when attention is restricted to symmetric equilibria.¹⁶ This section provides a condition on priors under which the present conditions are *strictly* weaker than the conditions for strategic proximity for general Bayesian games (and arbitrary equilibria) of Kajii and Morris [11], who show that a necessary condition for priors to be strategically close in general Bayesian games is that with high prior probability there is approximate common belief that conditional beliefs are close.

Before showing that the current conditions may be weaker than those for general Bayesian games for certain priors, it is worth asking why this may hold. Intuitively, the conditions for general Bayesian games for two priors to be close ensure that with high prior probability, there is no player who believes that with high conditional probability one of the other players believes that with high conditional probability . . . the conditional beliefs of some of his opponents are very different under the two priors. However, if priors assign positive probability only to interaction structures in which the actions of players who hold different conditional beliefs under the two priors (or the actions of players who believe that their opponents have different beliefs under the two priors, and so on) have limited impact on the payoffs of other players, then this condition may be too strong.¹⁷ To wit, even if a player of type t_1 believes that with high conditional probability the conditional beliefs of a player of type t_2 about some payoff-relevant event are very different under the two priors, the optimal strategy of the first player will be the same under the two priors if with high probability, a player of type t_2 has no effect on his payoffs. That is, only (higher-order) beliefs about neighbors matter, beliefs about arbitrary players who are not neighbors of a player do not affect equilibrium actions.

I now formalize this intuition. For $\delta > 0$, prior μ , type set $U \subseteq T_a$ and active type $t \in T_a$ such that $\mu(t) > 0$, say that there is a δ -path in neighbors from U to t (under μ), denoted

¹⁶The current conditions for strategic proximity for anonymous local games are necessary if arbitrary equilibria are considered, but need not be sufficient.

¹⁷In this section and the next, I follow standard terminology and speak of “players” rather than of agents or active players, as elsewhere in the paper. Also, the term “Bayesian game” refers to a Bayesian game with a finite player set.

$U \xrightarrow[\mu, \delta]{L} t$, if there is a sequence of players i_1, \dots, i_k such that each player i_m assigns conditional probability greater than δ to the event that one of his neighbors assigns conditional probability greater than δ to the event that... one of his neighbors has a type in U , i.e., there exists $k > 1$, $E_1, \dots, E_k \subseteq T_a$ with $E_1 = U$, $E_k = \{t\}$, such that for $m = 2, \dots, k$, there is $t_m \in E_m$ with $\mu(t_m) > 0$ and

$$\mu(\{(g, s) : \exists j \in N_1(g) \text{ s.t. } \tau_j(g, s) \in E_{m-1}\} \mid \{(g', s') : \tau_1(g', s') = t_m\}) > \delta, \quad (5)$$

where (with some abuse of notation) $N_1(g)$ and $\tau_1(g, s)$ are the set of neighbors and the type of the player associated with vertex 1 in g , respectively, and $\tau_j(g, s)$ is the type of an arbitrary neighbor of that player.

Similarly, there is a δ -path in players (under μ) from U to t (under μ), denoted $U \xrightarrow[\mu, \delta]{G} t$, if there exists $k > 1$, $E_1, \dots, E_k \subseteq T_a$ with $E_1 = U$, $E_k = \{t\}$, such that for $m = 2, \dots, k$, there is $t_m \in E_m$ with $\mu(t_m) > 0$ and

$$\mu(\{(g, s) : \exists j \neq 1 \text{ s.t. } \tau_j(g, s) \in E_{m-1}\} \mid \{(g', s') : \tau_1(g', s') = t_m\}) > \delta. \quad (6)$$

Condition (6) is almost identical to (5), with the important difference that a path in players only requires a sequence of (distinct) players, rather than a sequence of neighbors. The following result is thus immediate:

Lemma 4.5 *Let $\delta > 0$, $U \subseteq T_a$ and let $t \in T_a$ be such that $\mu(t) > 0$. If $U \xrightarrow[\mu, \delta]{L} t$, then $U \xrightarrow[\mu, \delta]{G} t$.*

A converse to this result clearly does not hold; for example, there is a δ -path in players between any pair of types for δ sufficiently small under the prior μ^* in Example 2.1, while there is no δ -path in neighbors from type t_ℓ to type t_m for $m < \ell$ for any $\delta > 0$.

How are these properties related to the conditions on priors to be strategically close? In Appendix B, I show that the conditions in Theorem 4.1 on players' posterior beliefs hold for a prior μ and a sequence $(\mu^\ell)_{\ell \in \mathbb{N}}$, while the corresponding condition for general Bayesian games is not satisfied if and only if:

- (a) with probability converging to one, there is *no* path in neighbors from the set of types that hold different conditional beliefs under μ and $(\mu^\ell)_{\ell \in \mathbb{N}}$ to a set of types of measure arbitrarily close to one; and
- (b) with probability bounded away from zero, there is a path in players from the set of types that have different conditional beliefs under μ and $(\mu^\ell)_{\ell \in \mathbb{N}}$ to a set of types of measure converging to one.

That is, suppose there is a prior μ and a sequence $(\mu^\ell)_{\ell \in \mathbb{N}}$ such that (a) and (b) hold, and the prior beliefs under μ^ℓ converge to those under μ (i.e., $d_0(\mu, \mu^\ell) \rightarrow 0$). Then there exists a

Bayesian game (μ, v) with anonymous local interactions such that μ^ℓ converges to μ in a strategic sense (i.e., $\chi^*(\mu, \mu^\ell; v, \varepsilon) \rightarrow 0$ for any $\varepsilon > 0$) even though the conditions for general Bayesian games are not satisfied. Example 2.1 provides an instance of this. As noted above, there is no path in neighbors from types with different conditional beliefs under μ and μ^ℓ to a set of types that has positive probability in the limit, and indeed players' equilibrium expected payoffs under $(\mu^\ell)_{\ell \in \mathbb{N}}$ come arbitrarily close to those under μ when ℓ grows large. On the other hand, there is no state in which there is approximate common belief that conditional beliefs are close, which means that the conditions for strategic closeness for general Bayesian games are not satisfied.

Intuitively, the above conditions on the prior effectively amount to restrictions on the possible payoff interdependencies between types, as some interactions have zero probability. If this is a reasonable assumption, and the interest is in the symmetric equilibria, then weaker conditions for strategic closeness suffice.

5 Network games

To illustrate the practical applicability of the present result, I now turn to the special case of network games, a class of games has which recently received considerable attention in the literature (e.g., [7, 8, 13]). A network game is a game in which players are located on a network, and play a fixed game with their neighbors. Players only know the number of neighbors they have, and are uncertain about the full network structure.

A central question in that literature is how the symmetric equilibria of a game change when players' beliefs about the network are varied. To illustrate how the present results can be used in that context, I will first derive the conditions for strategic proximity in network games. Since network games satisfy the anonymity and symmetry assumptions that define anonymous local games, Theorem 4.1 applies. However, because of the special structure of these games, weaker conditions suffice. Specifically, the private information of a player in a network game is the number of neighbors he has. Because the set of players is finite, this means that the type set is finite. In that case, uniform weak convergence of prior beliefs (i.e., $d_0(\mu, \mu^\ell) \rightarrow 0$) in fact implies uniform convergence of conditional beliefs. This follows from Lemma 5.1, which also provides a natural formulation of the condition on players' prior beliefs in the context of Bayesian network games:

Lemma 5.1 *Let $\mu \in \mathcal{M}$ and let $(\mu^\ell)_{\ell \in \mathbb{N}}$ be a sequence in \mathcal{M} . Assume that T_a is finite and that each type has positive probability under μ . Then, $\max_{F \subseteq T_a \times \Theta} |\mu(F) - \mu^\ell(F)| \rightarrow 0$ if and only if $\max_{t \in T_a, F \subseteq \Theta} |\mu(F | t) - \mu^\ell(F | t)| \rightarrow 0$, and $\max_{t \in T_a} |\mu(t) - \mu^\ell(t)| \rightarrow 0$.*

The proof can be found in the online appendix.

In turn, uniform convergence of conditional beliefs implies that $d_1(\mu, \mu^\ell)$ and $d_1(\mu^\ell, \mu)$ converge to zero as well. Consequently, weak convergence of prior beliefs is a necessary and sufficient condition for strategic convergence in Bayesian network games, which gives the following result:¹⁸

Corollary 5.2 *Consider a Bayesian network game (μ, v) such that the profile of payoff functions v is bounded, and let $(\mu^\ell)_{\ell \in \mathbb{N}}$ be a sequence of priors in \mathcal{M} . Fix $\varepsilon > 0$. Then*

$$\chi^*(\mu, \mu^\ell; v, \varepsilon) \rightarrow 0$$

if and only if for all $\zeta > 0$, there exists Λ_ζ such that for every $\ell > \Lambda_\zeta$,

$$\max_{t \in T_a, F \subseteq \Theta} |\mu(F | t) - \mu^\ell(F | t)| \leq \zeta,$$

and

$$\max_{t \in T_a} |\mu(t) - \mu^\ell(t)| \leq \zeta.$$

The intuition is the following. Suppose the prior probabilities of types and neighbor type profiles converge, i.e., $d_0(\mu, \mu^\ell) \rightarrow 0$. Because the type set is finite, the probability that a player has a given type is uniformly bounded away from zero. By Bayes' rule, the conditional belief of a player in an event that has low prior probability cannot be arbitrarily high. Hence, low-probability types cannot “infect” the rest of the type space, as in e.g. Example 2.2. That is, priors are insensitive to small probability events [11].

How can this result be used in the analysis of network games? Since very little is known about players' beliefs about their network, one might want to try to characterize the set of equilibrium outcomes in a network game for every possible prior, which is clearly infeasible. The present result identifies the regions in the space of all possible priors such that strategic predictions change continuously as beliefs are varied, as well as the regions where equilibrium behavior changes abruptly with small changes in beliefs. Corollary 5.2 thus says that two priors that induce similar distributions over the types of individual players, but differ in the correlation among types of neighbors (so that players' conditional beliefs over their neighbors' types will be different under the priors) can give very different equilibrium outcomes. This means that if one wants to explore the range of possible outcomes in a given network game, then it may not be enough to vary only the type distribution, as is done in much of the literature on network games¹⁹—also players' conditional beliefs need to be taken into account.

¹⁸This result of course applies to any setting with a finite type set. In fact, it applies more broadly to any setting where uniform convergence of prior probabilities implies uniform convergence of posteriors, as is for instance the case when types are independent [11].

¹⁹A notable exception is [7].

6 Concluding remarks

A common approach in the literature on the sensitivity of strategic predictions to higher-order uncertainty is to require some notion of continuity to hold uniformly over a wide range of games, for example the class of all Bayesian games with a countable type set. This paper departs from the literature in that respect by restricting attention to a subclass of games, characterized by certain restrictions on the payoff functions and the priors. The continuity of solution correspondences is then investigated relative to that class of games. Specifically, I derive necessary and sufficient conditions for priors to give similar equilibrium outcomes in anonymous local games such as network games and games with heterogeneous externalities. These conditions can be shown to be strictly weaker than the corresponding conditions for the class of all Bayesian games under certain conditions, so that it is possible to obtain stronger robustness results for this class of games.

It is likely that limiting the domain of games may also help weaken the conditions for robustness to higher-order uncertainty in other settings, such as the robustness of equilibria to small amounts of incomplete information (e.g., [10]). The challenge is to find reasonable restrictions on the class of games that allow for a significant weakening of the conditions for robustness for the larger class of games, preferably relaxing the condition that certain events need to be (approximate) common belief, in line with the goal set out by Wilson [26] of successively reducing the base of “common knowledge required to conduct useful analyses of practical problems.” The broader agenda of identifying interesting classes of games for which outcomes are more robust to higher-order uncertainty seems to be a promising one for future research.

Appendix A Proofs

A.1 Proof of Lemma 4.2

Lemma 4.2 uses Lemmas A.1–A.3.

Lemma A.1 *The set of types $C_\mu^p(U)$ is p -cohesive for any set $U \subseteq T_a$ of active types, i.e.,*

$$B_\mu^p(C_\mu^p(U)) = C_\mu^p(U).$$

Proof. By definition, $B_\mu^p(C_\mu^p(U)) \subseteq C_\mu^p(U)$. It remains to show that $B_\mu^p(C_\mu^p(U)) \supseteq C_\mu^p(U)$. Obviously, $([B_\mu^p]^\ell(U))_{\ell \in \mathbb{N}}$ is a weakly decreasing sequence, and, by definition, $\bigcap_{\ell \in \mathbb{N}} [B_\mu^p]^\ell(U) = C_\mu^p(U)$. Hence, using that the local p -belief operator is continuous (Lemma D.1 in the online appendix),

$$C_\mu^p(U) = \bigcap_{\ell \in \mathbb{N}} [B_\mu^p]^\ell(U) \subseteq \bigcap_{\ell \geq 2} [B_\mu^p]^\ell(U) = B_\mu^p\left(\bigcap_{\ell \in \mathbb{N}} [B_\mu^p]^\ell(U)\right) = B_\mu^p(C_\mu^p(U)),$$

and the result follows. \square

Lemma A.2 *Let $\mu \in \mathcal{M}$, and fix $\alpha, p \in [0, 1]$. For each set $U \subseteq T_a$ of active types, if the probability that an agent has a type in the set $C_\mu^p(U)$ is at least α , i.e., if*

$$\mu\left(\{(g, s) \in \mathcal{G} \times \mathcal{S} : \tau_1(g, s) \in C_\mu^p(U)\}\right) \geq \alpha,$$

then the probability that this agent and his neighbors all have a type in the set $C_\mu^p(U)$ is at least αp .

Proof. Recall that for any $U \subseteq T_a$, $\mu(\Xi(U) \mid t)$ denotes the conditional probability that the types of all neighbors of a given agent lie in U , given that the agent's type is t . By Lemma A.1, $C_\mu^p(U)$ is p -cohesive. Hence, for all $t \in C_\mu^p(U)$ such that $\mu(t) > 0$, $\mu(\Xi(C_\mu^p(U)) \mid t) \geq p$, so that

$$\begin{aligned} & \mu\left(\{(g, s) \in \mathcal{G} \times \mathcal{S} : \tau_1(g, s) \in C_\mu^p(U), \tau_j(g, s) \in C_\mu^p(U) \text{ for all } j \in N_1(g)\}\right) \\ &= \sum_{\substack{t' \in C_\mu^p(U): \\ \mu(t') > 0}} \frac{\mu(t') \cdot \mu\left(\{(g, s) \in \mathcal{G} \times \mathcal{S} : \tau_1(g, s) = t', \tau_j(g, s) \in C_\mu^p(U) \text{ for all } j \in N_1(g)\}\right)}{\mu\left(\{(g, s) \in \mathcal{G} \times \mathcal{S} : \tau_1(g, s) = t'\}\right)} \\ &\geq p \sum_{t' \in C_\mu^p(U)} \mu(t') \\ &\geq \alpha p, \end{aligned}$$

where, with some abuse of notation, $N_1(g)$ is the set of agents that are neighbors of the agent associated with vertex 1 in g . \square

Clearly, $C_\mu^p(U)$ can be replaced by any subset of T_a that is p -cohesive in the statement of Lemma A.2. It is presented in the present form for expositional reasons.

The next result builds on Lemma 1 of [22]:

Lemma A.3 *Let $\mu, \mu' \in \mathcal{M}$. If σ is a symmetric equilibrium of the anonymous local game (μ, v) and v is bounded by B , then there exists a symmetric $5\delta B$ -equilibrium σ' of the game (μ', v) such that $\sigma'_i(t) = \sigma_i(t)$ for any $i \in N_c$ and $t \in C_{\mu'}^{1-\delta}(T_{\mu, \mu'}^\delta)$.*

Proof. For ease of notation, define $Q := C_{\mu'}^{1-\delta}(T_{\mu, \mu'}^\delta)$, and note that Q is a subset of the set of active types T_a . Also, recall that $\Theta^{(k)}$ is the set of all neighbor type profiles of dimension k . Let $i \in N_c$. For $t \in Q$, set $\sigma'_i(t) = \sigma_i(t)$. For $t \notin Q$ such that $\mu'(t) > 0$, let $\sigma'_i(t)$ be such that $(\sigma'_j)_{j \in N_c}$ is a symmetric equilibrium of the reduced game in which each $j \in N_c$ with a type $t \in Q$ is committed to play $\sigma'_j(t) = \sigma_j(t)$. Such an equilibrium exists by Proposition 3.1. By construction, $\sigma'_j(t)$ is a best response to σ'_{-j} for any agent j of type $t \notin Q$. Hence, it remains to

show that $\sigma'_j(t)$ is a $5\delta B$ -best response to σ'_{-j} if j has type $t \in Q$. So let $t \in Q$, and recall that by definition, $\mu(t) > 0$ and $\mu'(t) > 0$ for $t \in Q$. By Lemma A.1,

$$\frac{\mu'(\{(g, s) \in \mathcal{G} \times \mathcal{S} : \tau_1(g, s) = t, \tau_j(g, s) \in Q \text{ for all } j \in N_1(g)\})}{\mu'(\{(g, s) \in \mathcal{G} \times \mathcal{S} : \tau_1(g, s) = t\})} \geq 1 - \delta, \quad (7)$$

where, again with some abuse of notation, $N_1(g)$ is the set of agents that are neighbors of the agent associated with vertex 1 in g . Moreover, by the definition of Q , for each subset of neighbor type profiles $F \subseteq \Theta$,

$$|\mu(F | t) - \mu'(F | t)| \leq \delta. \quad (8)$$

Consider an action $a \in A$ such that $\sigma_j(t)(a) > 0$, and let $b \in A$. Then,

$$\begin{aligned} |\varphi_j(a, \sigma'; t, \mu') - \varphi_j(b, \sigma'; t, \mu')| &\leq \sum_{k \in \mathbb{N}} \sum_{\theta \in \Theta^{(k)} \setminus Q^k} \mu'(\theta | t) |v_k(a, \sigma'_{(\theta)}, t, \theta) - v_k(b, \sigma'_{(\theta)}, t, \theta)| + \\ &\quad \sum_{k \in \mathbb{N}} \sum_{\theta \in Q^k \cap \Theta^{(k)}} \mu'(\theta | t) |v_k(a, \sigma'_{(\theta)}, t, \theta) - v_k(b, \sigma'_{(\theta)}, t, \theta)|, \end{aligned} \quad (9)$$

where $\mu'(\theta | t)$ denotes the probability (under the prior μ') that an agent's neighbor type profile is $\theta = (\theta_1, \dots, \theta_k)$ given that his type is t and where $\sigma'_{(\theta)}$ is the profile $(\sigma'_\ell(\theta_\ell))_{\ell=1, \dots, k}$ (recall that σ' is a symmetric strategy profile). The first sum in (9) can be evaluated directly. Using (7) and recalling that v is bounded by B ,

$$\sum_{k \in \mathbb{N}} \sum_{\theta \in \Theta^{(k)} \setminus Q^k} \mu'(\theta | t) |v_k(a, \sigma'_{(\theta)}, t, \theta) - v_k(b, \sigma'_{(\theta)}, t, \theta)| < \delta B. \quad (10)$$

To evaluate the second sum in (9), first note that if all neighbors of agent j have a type in set Q , they play according to the strategy profile σ . As σ is an equilibrium of (μ, v) , it follows from the triangle inequality that

$$\begin{aligned} \sum_{k \in \mathbb{N}} \sum_{\theta \in Q^k \cap \Theta^{(k)}} \mu(\theta | t) |v_k(a, \sigma_{(\theta)}, t, \theta) - v_k(b, \sigma_{(\theta)}, t, \theta)| &\leq \\ \sum_{k \in \mathbb{N}} \sum_{\theta \in \Theta^{(k)} \setminus Q^k} \mu(\theta | t) |v_k(a, \sigma_{(\theta)}, t, \theta) - v_k(b, \sigma_{(\theta)}, t, \theta)| \end{aligned} \quad (11)$$

Also, by (7) and (8), the conditional probability under μ that some neighbors of an agent do not have a type in Q , given that the agent's type is t , is at most 2δ . Combining this with (11) gives

$$\sum_{k \in \mathbb{N}} \sum_{\theta \in Q^k \cap \Theta^{(k)}} \mu(\theta | t) |v_k(a, \sigma_{(\theta)}, t, \theta) - v_k(b, \sigma_{(\theta)}, t, \theta)| \leq 2\delta B. \quad (12)$$

Let $P^k := \{\theta \in Q^k \cap \Theta^{(k)} : \mu'(\theta | t) - \mu(\theta | t) \geq 0\}$ be the set of neighbor type profiles θ in Q^k such that the conditional probability of θ under μ' is at least that under μ . Then, by (8) and

the triangle inequality,

$$\begin{aligned}
& \sum_{k \in \mathbb{N}} \sum_{\theta \in Q^k \cap \Theta^{(k)}} \left| (\mu'(\theta | t) - \mu(\theta | t)) (v_k(a, \sigma_{(\theta)}, t, \theta) - v_k(b, \sigma_{(\theta)}, t, \theta)) \right| = \\
& \quad \sum_{k \in \mathbb{N}} \sum_{\theta \in P^k \cap \Theta^{(k)}} (\mu'(\theta | t) - \mu(\theta | t)) |v_k(a, \sigma_{(\theta)}, t, \theta) - v_k(b, \sigma_{(\theta)}, t, \theta)| + \\
& \quad \sum_{k \in \mathbb{N}} \sum_{\theta \in (Q^k \cap \Theta^{(k)}) \setminus P^k} (\mu(\theta | t) - \mu'(\theta | t)) |v_k(a, \sigma_{(\theta)}, t, \theta) - v_k(b, \sigma_{(\theta)}, t, \theta)| \leq 2\delta B. \quad (13)
\end{aligned}$$

Combining (12) and (13) gives

$$\begin{aligned}
& \sum_{k \in \mathbb{N}} \sum_{\theta \in Q^k \cap \Theta^{(k)}} \mu'(\theta | t) |v_k(a, \sigma_{(\theta)}, t, \theta) - v_k(b, \sigma_{(\theta)}, t, \theta)| \leq \\
& \quad \sum_{k \in \mathbb{N}} \sum_{\theta \in Q^k \cap \Theta^{(k)}} \mu(\theta | t) |v_k(a, \sigma_{(\theta)}, t, \theta) - v_k(b, \sigma_{(\theta)}, t, \theta)| + \\
& \quad \sum_{k \in \mathbb{N}} \sum_{\theta \in Q^k \cap \Theta^{(k)}} |\mu'(\theta | t) - \mu(\theta | t)| |v_k(a, \sigma_{(\theta)}, t, \theta) - v_k(b, \sigma_{(\theta)}, t, \theta)| \leq 4\delta B. \quad (14)
\end{aligned}$$

Together with (9), (10), Eq. (14) yields

$$|\varphi_i(a, \sigma'; t, \mu') - \varphi_i(b, \sigma'; t, \mu')| \leq 5\delta B.$$

for any $i \in N_c$ and $t \in Q$. □

The proof of Lemma 4.2 now follows from these results. For ease of notation, define $Q := C_{\mu'}^{1-\delta}(T_{\mu, \mu'}^\delta)$, and note that $Q \subseteq T_a$. Since $d^*(\mu, \mu') \leq \delta$, the difference in probabilities assigned to the event that the type and neighbor type profile of an arbitrary agent belong to a subset F of $T_a \times \Theta$ under μ and μ' is at most δ , i.e.,

$$|\mu(F) - \mu'(F)| \leq \delta. \quad (15)$$

Moreover,

$$\mu'(\{(g, s) \in \mathcal{G} \times \mathcal{S} : \tau_1(g, s) \in Q\}) \geq 1 - \delta, \quad (16)$$

where $\tau_1(g, s)$ is the type of the agent associated with vertex 1 in g . Fix $i \in N_c$, and let σ be a symmetric equilibrium of (μ, v) . By Lemma A.3, there exists a symmetric $5\delta B$ -equilibrium σ' of (μ', v) such that $\sigma'_i(t) = \sigma_i(t)$ for any $i \in N_c$ and $t \in Q$. Hence, using (16) and Lemma A.2

(with $\alpha = p = 1 - \delta$),

$$\begin{aligned}
& |\Phi_i(\sigma'; \mu') - \Phi_i(\sigma; \mu')| \\
& \leq \sum_{t \in Q} \mu'(t) \sum_{\substack{k \in \mathbb{N}, \\ \theta \in Q^k \cap \Theta^{(k)}}} \mu'(\theta | t) \sum_{a \in A} |\sigma'_i(t)(a)v_k(a, \sigma'_{(\theta)}, t, \theta) - \sigma_i(t)(a)v_k(a, \sigma'_{(\theta)}, t, \theta)| + \\
& \sum_{t \in Q} \mu'(t) \sum_{\substack{k \in \mathbb{N}, \\ \theta \in \Theta^{(k)} \setminus Q^k}} \mu'(\theta | t) \sum_{a \in A} |\sigma'_i(t)(a)v_k(a, \sigma'_{(\theta)}, t, \theta) - \sigma_i(t)(a)v_k(a, \sigma'_{(\theta)}, t, \theta)| + \\
& \sum_{\substack{t \in T_a \setminus Q: \\ \mu'(t) > 0}} \mu'(t) \sum_{\substack{k \in \mathbb{N}, \\ \theta \in \Theta^{(k)}}} \mu'(\theta | t) \sum_{a \in A} |\sigma'_i(t)(a)v_k(a, \sigma'_{(\theta)}, t, \theta) - \sigma_i(t)(a)v_k(a, \sigma'_{(\theta)}, t, \theta)| \\
& < 0 + (1 - (1 - \delta)^2)B \\
& = (2 - \delta)\delta B,
\end{aligned} \tag{17}$$

where $\sigma'_{(\theta)} = (\sigma'_\ell(\theta_\ell))_{\ell=1, \dots, k}$ for $\theta = (\theta_1, \dots, \theta_k) \in \Theta^{(k)}$ (recall that σ' is a symmetric strategy profile, so that the agents' identities are irrelevant). Let $P := \{t \in T_a : \mu'(t) - \mu(t) \geq 0\}$. Then,

$$\begin{aligned}
|\Phi_i(\sigma; \mu') - \Phi_i(\sigma; \mu)| & \leq \sum_{t \in P} \sum_{\substack{k \in \mathbb{N}, \\ \theta \in \Theta^{(k)}}} (\mu'(t) - \mu(t)) \sum_{a \in A} \sigma_i(t)(a) |v_k(a, \sigma_{(\theta)}, t, \theta)| + \\
& \sum_{t \in T_a \setminus P} \sum_{\substack{k \in \mathbb{N}, \\ \theta \in \Theta^{(k)}}} (\mu(t) - \mu'(t)) \sum_{a \in A} \sigma_i(t)(a) |v_k(a, \sigma_{(\theta)}, t, \theta)|, \\
& \leq 2\delta B.
\end{aligned} \tag{18}$$

Combining (17) and (18) gives the desired result. \square

A.2 Proof of Lemma 4.3

If $d_0(\mu, \mu') > \delta$, there exists $F \subseteq T_a \times \Theta$ such that $|\mu(F) - \mu'(F)| > \delta$, where $\mu(F)$ and $\mu'(F)$ denote the prior probability under μ and μ' that an arbitrary agent has a type and neighbor type profile in F , respectively. For each $t \in T_a, k > 0, b \in A, a^k \in A^k$ and $\theta \in \Theta^{(k)}$, let

$$v_k(b, a^k, t, \theta) = \begin{cases} 1 & \text{if } (t, \theta) \in F, \\ 0 & \text{otherwise.} \end{cases}$$

It is easy to see that for all $i \in N_c$,

$$|\Phi_i(\sigma; \mu) - \Phi_i(\sigma'; \mu')| > \delta$$

for *any* strategy profiles σ, σ' . \square

Appendix B Conditions for Bayesian games

This appendix relates the conditions for strategic proximity in anonymous local games and general Bayesian games to the probability that there are paths in neighbors and paths in players, respectively. Denote the fixed (finite) set of players by I .

I show the following result. Consider a prior μ and a sequence $(\mu^\ell)_{\ell \in \mathbb{N}}$. The conditions $d_1(\mu, \mu^\ell) \rightarrow 0$ and $d_1(\mu^\ell, \mu) \rightarrow 0$ on players' conditional beliefs hold, while the corresponding condition for general Bayesian games (specified below) is not satisfied if and only if for all $\zeta > 0$, there exists Λ_ζ such that for all $\ell > \Lambda_\zeta$,

$$\inf \left\{ \delta : \mu \left(\left\{ (g, s) : \tau_1(g, s) \notin T_{\mu, \mu^\ell}^\delta \text{ or } \exists t \notin T_{\mu, \mu^\ell}^\delta \text{ s.t. } t \xrightarrow[\mu, \delta]{L} \tau_1(g, s) \right\} \leq \delta \right) \leq \zeta, \text{ and} \right. \\ \left. \inf \left\{ \delta : \mu^\ell \left(\left\{ (g, s) : \tau_1(g, s) \notin T_{\mu, \mu^\ell}^\delta \text{ or } \exists t \notin T_{\mu, \mu^\ell}^\delta \text{ s.t. } t \xrightarrow[\mu^\ell, \delta]{L} \tau_1(g, s) \right\} \leq \delta \right) \leq \zeta \right. \right.$$

and there exists $c > 0$ such that

$$\inf \left\{ \delta : \mu \left(\left\{ (g, s) : \exists j \in I \text{ s.t. } \tau_j(g, s) \notin T_{\mu, \mu^\ell}^\delta \text{ or } \exists t \notin T_{\mu, \mu^\ell}^\delta \text{ s.t. } t \xrightarrow[\mu, \delta]{G} \tau_j(g, s) \right\} \leq \delta \right) > c, \text{ or} \right. \\ \left. \inf \left\{ \delta : \mu^\ell \left(\left\{ (g, s) : \exists j \in I \text{ s.t. } \tau_j(g, s) \notin T_{\mu, \mu^\ell}^\delta \text{ or } \exists t \notin T_{\mu, \mu^\ell}^\delta \text{ s.t. } t \xrightarrow[\mu^\ell, \delta]{G} \tau_j(g, s) \right\} \leq \delta \right) > c \right. \right.$$

for all ℓ sufficiently large. This is done by showing that the prior probability that a player has a type in the cohesive subset of types that have similar conditional beliefs under two priors is equal to the prior probability that there is no path in neighbors from a type that has different conditional beliefs under these priors to that player's type. Similarly, the prior probability that there is approximate common belief that the conditional beliefs of all players are close under two priors equals the prior probability that there is no path in players from the types that hold different conditional beliefs to the players' (realized) types.

I first recall some notation from [11]. For any set of states $F_G \subseteq \mathcal{G} \times \mathcal{S}$ and $p \in [0, 1]$, let

$$B_{\mu, G}^p(F_G) := \{(g, s) \in \mathcal{G} \times \mathcal{S} : \forall i \in I, \mu(\tau_i(g, s)) > 0 \implies \mu(F_G \mid \tau_i(g, s)) \geq p\},$$

where $\mu(\tau_i(g, s))$ is the prior probability that player i has type $\tau_i(g, s)$ and $\mu(F_G \mid \tau_i(g, s))$ is i 's posterior belief that the true state lies in the set F_G . The operator $B_{\mu, G}^p$ is the p -belief operator introduced by Monderer and Samet [15]. Iterating the belief operator gives approximate common belief in the event F_G :

$$C_{\mu, G}^p(F_G) = \bigcap_{\ell \in \mathbb{N}} [B_{\mu, G}^p]^\ell(F_G).$$

Finally, for priors μ, μ' , let

$$\mathcal{A}_{\mu, \mu'}^\delta := \left\{ (g, s) \in \mathcal{G} \times \mathcal{S} : \forall i \in I, \mu(\tau_i(g, s)), \mu'(\tau_i(g, s)) > 0, \text{ and} \right. \\ \left. \sup_{F_G \subseteq \mathcal{G} \times \mathcal{S}} \left| \mu(F_G \mid \tau_i(g, s)) - \mu'(F_G \mid \tau_i(g, s)) \right| \leq \delta \right\}$$

be the set of states such that the conditional beliefs of all agents are close under μ and μ' . Kajii and Morris [11] show that the following condition is necessary for strategic convergence of a sequence $(\mu^\ell)_{\ell \in \mathbb{N}}$ to a prior μ on the class of Bayesian games: For every $\zeta > 0$, there exists Λ_ζ such that for all $\ell > \Lambda_\zeta$,

$$\inf \left\{ \delta : \mu \left(C_{\mu, G}^{1-\delta}(\mathcal{A}_{\mu, \mu^\ell}^\delta) \right) \geq 1 - \delta \right\} \leq \zeta,$$

and

$$\inf \left\{ \delta : \mu^\ell \left(C_{\mu^\ell, G}^{1-\delta}(\mathcal{A}_{\mu, \mu^\ell}^\delta) \right) \geq 1 - \delta \right\} \leq \zeta.$$

To explore the conditions on priors for which the above condition is weaker than the conditions in Theorem 4.1 on players' conditional beliefs (i.e., $d_1(\mu, \mu^\ell) \rightarrow 0$ and $d_1(\mu^\ell, \mu) \rightarrow 0$), the following two results relate the different conditions for strategic proximity to the presence of paths in neighbors and paths in players, respectively:

Proposition B.1 *Fix $\delta > 0$ and let μ, ν be priors on $\mathcal{G} \times \mathcal{S}$. Then,*

$$\left\{ (g, s) \in \mathcal{G} \times \mathcal{S} : \tau_1(g, s) \notin T_{\mu, \nu}^\delta \text{ or } \mu(\tau_1(g, s)) > 0 \text{ and } T_a \setminus T_{\mu, \nu}^\delta \xrightarrow[\mu, \delta]{L} \tau_1(g, s) \right\} = \\ \left\{ (g, s) \in \mathcal{G} \times \mathcal{S} : \tau_1(g, s) \notin C_\mu^{1-\delta}(T_{\mu, \nu}^\delta) \right\}$$

Proof. Fix a state $(g, s) \in \mathcal{G} \times \mathcal{S}$. If the type of a player associated with vertex 1 in g is not in $T_{\mu, \nu}^\delta$, i.e., $\tau_1(g, s) \notin T_{\mu, \nu}^\delta$, then $\tau_1(g, s) \notin C_\mu^{1-\delta}(T_{\mu, \nu}^\delta)$. So suppose $\mu(\tau_1(g, s)) > 0$ and $T_a \setminus T_{\mu, \nu}^\delta \xrightarrow[\mu, \delta]{L} \tau_1(g, s)$. Then there exist $k, E_1, \dots, E_k \subseteq T_a$ with $E_1 := T_a \setminus T_{\mu, \nu}^\delta$, $E_k := \{\tau_1(g, s)\}$, and for all $m \geq 2$, $t_m \in E_m$, $\mu(t_m) > 0$ and

$$\mu(\{(g', s') : \exists j \in N_1(g') \text{ s.t. } \tau_j(g', s') \in E_{m-1}\} \mid \{(g'', s'') : \tau_1(g'', s'') = t_m\}) > \delta.$$

Using that $E_1 = T_a \setminus T_{\mu, \nu}^\delta$, it follows that for each $m \geq 2$, $t_m \in E_m$,

$$\mu(\{(g', s') : \forall j \in N_1(g'), \tau_j(g', s') \in [B_\mu^{1-\delta}]^{m-2}(T_{\mu, \nu}^\delta)\} \mid \{(g'', s'') : \tau_1(g'', s'') = t_m\}) < 1 - \delta,$$

so that $t_m \notin [B_\mu^{1-\delta}]^{m-1}(T_{\mu, \nu}^\delta)$. This implies that $\tau_1(g, s) \notin C_\mu^{1-\delta}(T_{\mu, \nu}^\delta)$.

Conversely, suppose $\tau_1(g, s) \notin C_\mu^{1-\delta}(T_{\mu, \nu}^\delta)$ for some fixed $(g, s) \in \mathcal{G} \times \mathcal{S}$. If $\tau_1(g, s) \notin T_{\mu, \nu}^\delta$, the proof is complete. So suppose that $\mu(\tau_1(g, s)) > 0$ and $\tau_1(g, s) \in T_{\mu, \nu}^\delta$. Then there exists a smallest $k \geq 1$ such that $\tau_1(g, s) \notin [B_\mu^{1-\delta}]^k(T_{\mu, \nu}^\delta)$. Let $U_1 := T_a \setminus T_{\mu, \nu}^\delta$. Note that U_1 is nonempty: If $U_1 = \emptyset$, then $C_\mu^{1-\delta}(T_{\mu, \nu}^\delta) = T_{\mu, \nu}^\delta$, contradicting that $\tau_1(g, s) \in T_{\mu, \nu}^\delta \setminus C_\mu^{1-\delta}(T_{\mu, \nu}^\delta)$. For $m = 2, \dots, k$, define

$$U_m := \{t \in T_a \setminus [B_\mu^{1-\delta}]^{m-1}(T_{\mu, \nu}^\delta) : \mu(t) > 0\},$$

and note that $U_m \subseteq U_{m+1}$ for all m . Finally, let $U_{k+1} := \{\tau_j(g, s)\}$. It follows that for $m = 2, \dots, k$ and $t_m \in U_m$, $\mu(t_m) > 0$ and $t_m \notin [B_\mu^{1-\delta}]^{m-1}(T_{\mu,\nu}^\delta)$, so that

$$\begin{aligned} & \mu(\{(g', s') : \exists j \in N_1(g') \text{ s.t. } \tau_j(g', s') \in U_{m-1}\} \mid \{(g'', s'') : \tau_1(g'', s'') = t_m\}) \\ &= \mu(\{(g', s') : \exists j \in N_1(g') \text{ s.t. } \tau_j(g', s') \notin [B_\mu^{1-\delta}]^{m-2}(T_{\mu,\nu}^\delta)\} \mid \{(g'', s'') : \tau_1(g'', s'') = t_m\}) \\ &> \delta. \end{aligned}$$

Finally, as $\tau_1(g, s) \notin [B_\mu^{1-\delta}]^k(T_{\mu,\nu}^\delta)$,

$$\begin{aligned} & \mu(\{(g', s') : \exists j \in N_1(g') \text{ s.t. } \tau_j(g', s') \in U_k \mid \{(g'', s'') : \tau_1(g'', s'') = \tau_1(g, s)\}\}) \\ &= \mu(\{(g', s') : \exists j \in N_1(g') \text{ s.t. } \tau_j(g', s') \notin [B_\mu^{1-\delta}]^{k-1}(T_{\mu,\nu}^\delta) \mid \{(g'', s'') : \tau_1(g'', s'') = \tau_1(g, s)\}\}) \\ &> \delta, \end{aligned}$$

so that there is a δ -path in neighbors from $T_a \setminus T_{\mu,\nu}^\delta$ to $\tau_1(g, s)$. \square

Proposition B.2 *Fix $\delta > 0$ and let μ, ν be priors on $\mathcal{G} \times \mathcal{S}$. Then,*

$$\left\{ (g, s) \in \mathcal{G} \times \mathcal{S} : \exists j \in I \text{ s.t. } \tau_j(g, s) \notin T_{\mu,\nu}^\delta \text{ or } T_a \setminus T_{\mu,\nu}^\delta \xrightarrow[\mu,\delta]{G} \tau_j(g, s) \right\} = (\mathcal{G} \times \mathcal{S}) \setminus C_{\mu,G}^{1-\delta}(\mathcal{A}_{\mu,\nu}^\delta).$$

Proof. Note that if (g, s) is such that there is $j \in I$ with $\tau_j(g, s) \notin T_{\mu,\nu}^\delta$, then $(g, s) \notin \mathcal{A}_{\mu,\nu}^\delta$. Using this, the proof is analogous to the proof of Proposition B.1. The details are omitted. \square

Noting that if there is a player in $j \in I$ such that $\tau_j(g, s) \notin T_{\mu,\nu}^\delta$, then $(g, s) \notin \mathcal{A}_{\mu,\nu}^\delta$, it follows from Propositions B.1 and B.2 that

$$\{(g, s) : \tau_1(g, s) \notin C_\mu^{1-\delta}(T_{\mu,\nu}^\delta)\} \subseteq (\mathcal{G} \times \mathcal{S}) \setminus C_{\mu,G}^{1-\delta}(\mathcal{A}_{\mu,\mu^\ell}^\delta).$$

Lemma 4.5 then gives the desired result.

References

- [1] E. J. Balder and A. Rustichini. An equilibrium result for games with private information and infinitely many players. *J. Econ. Theory*, 62:385–393, 1994.
- [2] S. Bernstein and E. Winter. Contracting with heterogeneous externalities. Working paper, Hebrew University, Center for the Study of Rationality, 2009.
- [3] H. Carlsson and E. van Damme. Global games and equilibrium selection. *Econometrica*, 61:989–1018, 1993.

- [4] Y.-C. Chen, A. Di Tillio, E. Faingold, and S. Xiong. Uniform topologies on types. *Theoretical Economics*, 5:445–478, 2010.
- [5] M. H. DeGroot. *Optimal Statistical Decisions*. Mcgraw-Hill College, 1970.
- [6] E. Dekel, D. Fudenberg, and S. Morris. Topologies on types. *Theoretical Economics*, 1:275–309, 2006.
- [7] A. Galeotti, S. Goyal, M. O. Jackson, F. Vega-Redondo, and L. Yariv. Network games. *Rev. Econ. Stud.*, 77:218–244, 2010.
- [8] M. O. Jackson and L. Yariv. Diffusion of behavior and equilibrium properties in network games. *Amer. Econ. Rev. (Papers and Proceedings)*, 97(2):92–98, 2007.
- [9] Ph. Jehiel and B. Moldovanu. Strategic nonparticipation. *RAND J. Econ.*, 27:84–98, 1996.
- [10] A. Kajii and S. Morris. The robustness of equilibria to incomplete information. *Econometrica*, 65:1283–1309, 1997.
- [11] A. Kajii and S. Morris. Payoff continuity in incomplete information games. *J. Econ. Theory*, 82:267–276, 1998.
- [12] E. Kalai. Large robust games. *Econometrica*, 72:1631–1666, 2004.
- [13] D. López-Pintado. Diffusion in complex social networks. *Games Econ. Behav.*, 62:573–590, 2008.
- [14] I. Milchtaich. Random-player games. *Games Econ. Behav.*, 47:353–388, 2004.
- [15] D. Monderer and D. Samet. Approximating common knowledge with common beliefs. *Games Econ. Behav.*, 1:170–190, 1989.
- [16] D. Monderer and D. Samet. Proximity of information in games with incomplete information. *Math. Oper. Res.*, 3:707–725, 1996.
- [17] S. Morris. Interaction games: A unified analysis of incomplete information, local interaction and random matching games. Working paper 97-08-072E, Santa Fe Institute, 1997.
- [18] S. Morris. Contagion. *Rev. Econ. Stud.*, 67:57–78, 2000.
- [19] S. Morris and H. S. Shin. Common belief foundations of global games. Working paper, Princeton, 2007.

- [20] R. B. Myerson. Population uncertainty and Poisson games. *Int. J. Game Theory*, 27:375–392, 1998.
- [21] B. Peleg. Equilibrium points for games with infinitely many players. *J. London Math. Soc.*, 44:292–294, 1969.
- [22] C. G. Rothschild. Payoff continuity in incomplete information games: A comment. *J. Econ. Theory*, 120:270–274, 2005.
- [23] A. Rubinstein. The electronic mail game: Strategic behavior under “almost common knowledge”. *Amer. Econ. Rev.*, 79:385–391, 1989.
- [24] J. L. Weinstein and M. Yildiz. Impact of higher-order uncertainty. *Games Econ. Behav.*, 60:200–212, 2007.
- [25] J. L. Weinstein and M. Yildiz. A structure theorem for rationalizability with application to robust predictions of refinements. *Econometrica*, 75:365–400, 2007.
- [26] R. B. Wilson. Game-theoretic analyses of trading processes. In T. Bewley, editor, *Advances in Economic Theory: Fifth World Congress*, pages 33–70. Cambridge Univ. Press, Cambridge, UK, 1987.