

Distance-Based Equilibria in Normal-Form Games

Erman Acar

Vrije Universiteit Amsterdam
Amsterdam, The Netherlands
erman.acar@vu.nl

Reshef Meir

Technion - Israel Institute of Technology
Haifa, Israel
reshefm@ie.technion.ac.il

Abstract

We propose a simple uncertainty modification for the agent model in normal-form games; at any given strategy profile, the agent can access only a set of “possible profiles” that are within a certain distance from the actual action profile. We investigate the various instantiations in which the agent chooses her strategy using well-known rationales e.g., considering the worst case, or trying to minimize the regret, to cope with such uncertainty. Any such modification in the behavioral model naturally induces a corresponding notion of equilibrium; a *distance-based equilibrium*. We characterize the relationships between the various equilibria, and also their connections to well-known existing solution concepts such as Trembling-hand perfection. Furthermore, we deliver existence results, and show that for some class of games, such solution concepts can actually lead to better outcomes.

Introduction

Decision making under uncertainty is a key issue both in game theory and in artificial intelligence. Whereas models of strict uncertainty, or absence of unique priors are common at the outset of those fields (Gilboa and Schmeidler 1989; Dow and Werlang 1994; Halpern 2017; Gilboa, Postlewaite, and Schmeidler 2008; Potyka et al. 2016), probabilistic and Bayesian models still have become dominant, albeit due to different reasons. In AI, the use of probabilities often leads to better performance in a wide variety of tasks (e.g. Bayesian networks for debugging and information retrieval (Heckerman, Mamdani, and Wellman 1995), Monte-Carlo methods for robot localization (Thrun et al. 2001), and many more). In game theory, probabilities are used first and foremost because they allow for clean modeling, and in particular the use of von Neumann-Morgenstern utilities with all the rich theory that they support.

Several solution concepts suggested in the behavioral game theory literature tackled the problem of imperfect rationality (e.g., Cognitive hierarchy (Camerer, Ho, and Chong 2004), Quantal response (McKelvey and Palfrey 1995), Trembling-hand perfect equilibria (Selten 1975) and others), yet many of these still assume that agents optimize or approximate their expected utility over some distribution, which is not very cognitively plausible. For exam-

ple, following a Trembling-hand perfect strategy, which appears to be the foremost central notion of equilibrium refinements, can be claimed to be a “super-rational” behavior rather than bounded rational (Aumann 1997). Recently, Trembling-hand perfect equilibrium regained strong attention in machine learning research as well (Farina et al. 2018; Farina, Gatti, and Sandholm 2018).

In this work, we suggest a model for an agent inspired by early work in AI on uncertainty and reasoning (see the work by Halpern (Halpern 2017)), and some recent works on specific games (such as voting (Conitzer, Walsh, and Xia 2011; Meir, Lev, and Rosenschein 2014; Lev et al. 2019) and routing (Meir and Parkes 2015)), and look at its very foundations; normal-form games, from the lenses of this model. Our model is *distance-based*, in the sense that at any given action profile, a set of “possible profiles” (that are close to the actual profile) is constructed w.r.t. a metric, without assigning them any particular single probability. Then, the agent determines her action using one of the many available rationales for decision making under strict uncertainty, e.g., considering the worst case (Wald 1939) or trying to minimize regret (Savage 1951; Hyafil and Boutilier 2004).

The intuition of such setting is due to imprecision caused by limitations in observations (of the agent), and the “closeness” of signals (imposed by the environment) which causes perceptual indistinguishability that comes with it.¹ Intuitively, to capture such notions formally, one can employ a distance-based model. Note also that such model can be considered as a bounded rationality model, as it limits the agent’s reasoning about the other agents’ strategies.

It is flexible in the sense that potentially different distance metrics and decision rules can be plugged into the model, to fit specific games or types of behavior. Once we fix our behavioral model though, it naturally induces a notion of *equilibrium*, which is an action profile where no agent is inclined to change her action.

In *biased games* (Caragiannis, Kurokawa, and Procaccia 2014), a subclass of *penalty games* (Deligkas, Fearnley, and Spirakis 2016)), players are equipped with a non-linear utility function; this is due to an additional bias term (or *penalty*)

¹For instance, see *random error* (Cohen 1998). Note that these discussions also took place in philosophy and epistemology e.g., “distant trees” argument (Williamson 1992).

occurring in the utility function. The bias term itself is a real-valued function defined on the distance (via an L_p norm) between the played strategy and a *base strategy* (a particular strategy e.g., represents a social norm). These particular features and the results follow, stand orthogonal to our work.

Some other related works (and the references therein) that are worth mentioning: The authors in (Marinacci 2000) makes use of Choquet expected utility model based on non-additive probabilities (Gilboa and Schmeidler 1989). Their pessimistic/optimistic choices (Marinacci 2000) share similar intuition with our worst-case/best-case responses. The notion of local rationalizability by K. Apt (Apt 2007) seems related to our *local-best response*; yet there it is enough for a strategy to be a best-response to a single strategy in beliefs, whereas in our case it has to be the one optimal w.r.t the whole belief set. Aghassi and Bertsimas in their work (Aghassi and Bertsimas 2006) uses robust optimization (hence worst-case scenario) to model uncertainty in payoffs. None of those works, however, uses the distance as the basic machinery, and different technical subtleties and challenges apply.

Contribution and paper structure After giving basics and the familiar equilibrium notions in the next section, we introduce our model and explore the interlinks between different variations of it. We also explore its relation to other major refinements, among others the aforementioned notion of Trembling-hand perfection. Most of our results are not metric-specific, yet in examples and some results, due to its wide-spread use and intuition, we adopt Euclidean metric. Then we demonstrate how these solution concepts apply to several common games of interest, and provide with existence results for our notion. To underline its benefit, we introduce a class of games such that these notions potentially guarantee better outcomes. And very much in connection with that, as our final contribution, we provide a result which gives a price of anarchy bound in terms of our notion. Conclusion and future work closes the paper.

Preliminaries and Notation

We define n -player *normal-form game* $\mathcal{G} = (N, A, u)$, where N is a finite set of n players, indexed by i ; $A = A_1 \times \dots \times A_n$, where A_i is a finite set of *actions* (or *pure strategies*) available to player i .

Each vector $a = (a_1, \dots, a_n) \in A$ is called an action profile; $u = (u_1, \dots, u_n)$ where $u_i : A \rightarrow \mathbb{R}$ is a real-valued *utility function* (or *payoff function*) for player i . A mixed strategy π_i for player i is a probability distribution over the set of available actions A_i for player i . Further, we denote the set of *mixed-strategies* for player i by Π_i , which is the set of all probability distributions over the set A_i of actions for player i . The set of mixed-strategy profiles is simply the Cartesian product of the individual mixed-strategy sets i.e., $\Pi = \Pi_1 \times \dots \times \Pi_n$. We denote a (mixed-strategy) profile by $\pi \in \Pi$. Further, for a player i , we denote the probability that an action a_i is played under mixed strategy π_i , by $\pi_i(a_i)$. The *support* of a mixed strategy π_i for a player i is the set of pure strategies $\{a_i | \pi_i(a_i) > 0\}$. For a player i , a

mixed-strategy π_i is *totally* (or *completely*) *mixed* if its support subsumes A_i . A strategy profile π is totally mixed if its every component is totally mixed.

For simplicity, we overload the function symbol u_i to define the (expected) utility u_i of a strategy profile π for player i in a normal-form game as $u_i(\pi) = \sum_{a \in A} u_i(a) \prod_{j \in N} \pi_j(a_j)$.² A (mixed) strategy π_i is a *best response* to π_{-i} if $u_i(\pi_i, \pi_{-i}) \geq u_i(\pi'_i, \pi_{-i})$ for every $\pi'_i \in \Pi_i$. A (mixed) strategy profile is a *Nash equilibrium* (*MN*) if for every player $i \in N$, π_i is a best-response to π_{-i} . A pure strategy Nash equilibrium (*PN*) is a *MN* where every player's strategy has a support of cardinality 1. A *totally mixed Nash equilibrium* is denoted by *TMN*. Given a profile π , *Social Welfare* $SW(\pi) = \sum_{i \in N} u_i(\pi)$. And finally, *Price of Anarchy* PoA for a game is defined as the ratio of the maximum social welfare (numerator) to the minimum social welfare in an equilibrium (denominator).

Equilibria with Mistakes and Imprecision

We mention definitions of several well-known equilibrium concepts involved with *slight mistakes* or *imprecision* of agents, from the literature. These are *Trembling-Hand Perfect Equilibrium* (*T*) from R. Selten's seminal work (Selten 1975), its stronger version, *Truly Perfect Equilibrium* (*TP*) (Kohlberg 1981), and *Robust equilibrium* (*R*) (Messner and Polborn 2005). These concepts are of particular importance since we shall reveal their connections to the distance-based equilibrium concepts, introduced later in the next section.

Definition 1 (Trembling-Hand Perfect Equilibrium (Selten 1975)). *Given a finite game \mathcal{G} , a mixed strategy profile π is Trembling-hand perfect equilibrium if there is a sequence $\{\pi^k\}_{k=0}^{\infty}$ of totally mixed strategy profiles which converges to π such that for each agent $i \in N$, π_i is a best response to π_{-i}^k for all k .*

Selten's notion of Trembling-hand perfect equilibrium is based on the notion of best response which is robust against minimal mistakes (hence the term *trembling hand*) of opponents, formalized by a sequence of profiles converging to the equilibrium. As it was shown by Selten (Selten 1975), every finite game has a *T*-equilibrium. Note that the notion does not demand a best-response to every such sequence but rather only one. We shall later show that this very notion is entangled to our notions of distance-based equilibria. So is the Truly Trembling-hand perfect equilibrium, a stronger variant, as we mention next.³

Definition 2 (Truly perfect equilibrium (Kohlberg 1981)). *Given a finite game \mathcal{G} , a mixed strategy profile π is truly perfect equilibrium if for each sequence $\{\pi^k\}_{k=0}^{\infty}$ of totally mixed strategy profiles which converge to π , there is a K such that π_i is a best response to π_{-i}^k for all $k \geq K$.*

It is easy to see that this notion demands a lot by requiring each action to be a best response in every sequence of profile converging to the equilibrium. There is a cost for this

²Note that large \prod in this expression stands for product (instead of Π , the set of mixed profiles).

³Another similar variation (Okada 1981) that aims to strengthen Selten's Trembling-Hand is given by Okada.

demand, that is, TP does not always exist (Kohlberg 1981) (see also Chapter 11 of (Fudenberg and Tirole 1991)).

We also define an even stronger variation; namely, *Strict-TP* ($s-TP$).

Definition 3 (Strict- TP). *A strategy profile is a Strict- TP if it is defined as in Def. 2 except that every π_i is strictly better than any other response to π_{-i}^k .*

We also provide a slightly stronger extension of a Selten's original Trembling-hand perfect equilibrium; namely *strict- T* ($s-T$):

Definition 4 (Strict- T). *A strategy profile is a Strict- T if it is defined as in Def. 1 except that for every $i \in N$, there is an infinite subsequence of $\{\pi^k\}_{k=0}^\infty$ where π_i is a strict best-response to π_{-i}^k .*

Note that a strict- T has to be a pure Nash equilibrium, since a strict best-response cannot be mixed.

The following notion of Robust equilibrium is an adaptation from robust political equilibrium (Messner and Polborn 2005).⁴ Intuitively, an equilibrium is ϵ -Robust, if each player would like to keep her action, even if there is a small chance that other players deviate.

Definition 5 (Robust Equilibrium (Messner and Polborn 2005)). *A mixed profile π is an ϵ -noisy variant of a pure profile a , if for all $j \in N$, $\pi_j(a_j) > 1 - \epsilon$ where $\epsilon > 0$.*

- Given a pure strategy profile a , player i , and $\epsilon > 0$, action b_i is an ϵ -Robust response if b_i is a best response to any ϵ -noisy variant of a_{-i} .
- A pure strategy profile a is an ϵ -Robust equilibrium if every a_i is an ϵ -Robust response to a_{-i} .

Next, we provide a link between those two concepts.

Proposition 1. *If a is an ϵ -Robust equilibrium for some $\epsilon > 0$, then a is a TP .*

Proof. Let a be some ϵ -Robust equilibrium for some $\epsilon > 0$, and consider a sequence $\{\pi^k\}_{k=0}^\infty$ converging to a . Thus $\pi_i^k(a_i) \rightarrow 1$ for all $i \in N$. In particular, there is some K_i such that for all $k > K_i$, $\pi_i^k(a_i) > 1 - \epsilon$. Let $K := \max_i K_i$. Then for all $k > K$, and for all $j \in N$, we have that $\pi_j^k(a_j) > 1 - \epsilon$, i.e. π^k is an ϵ -noisy variant of a , and thus a_{-i} is a best response to π_{-i}^k . \square

Distance-based Equilibria

In the following, we define the central notions of the paper.

Distance-based uncertainty

For every agent $i \in N$, let $r_i \in \mathbb{R}^+$ be the associated *ignorance factor* formalizing the intuition: Greater the ignorance factor, more ignorant/cautious the agent (about the mixed strategies of other agents).

⁴(Messner and Polborn 2005) deal with both coalitional stability and noisy actions, assuming that each player "fails" to play with some small probability. We only focus on the latter part. Our definition of Robust equilibrium is based on their informal description and motivation.

Given a mixed strategy profile $\pi = \langle \pi_i, \pi_{-i} \rangle$, $\mathcal{B}_i(\pi, r_i) := \{\pi'_{-i} \mid d(\pi_{-i}, \pi'_{-i}) \leq r_i\}$ is the set of *possible response profiles* of others that the agent i is considering, equipped with a metric d which is assumed to have the axioms of non-negativity i.e., $d(x, y) \geq 0$; identity of indiscernibles i.e., $d(x, y) = 0 \iff x = y$; symmetry i.e., $d(x, y) = d(y, x)$; and triangle inequality i.e., $d(x, z) \leq d(x, y) + d(y, z)$ for all $x, y, z \in \Pi$. Note that although our results comply with any compact space with a metric which fulfills those axioms, in the examples throughout the paper we work with Euclidean metric for convenience (due to its wide-spread use). Often, we will use the shorthand notation $\mathcal{B}_i(\pi)$ (\mathcal{B} for ball) whenever r_i is clear from the context.

Intuitively, $\mathcal{B}_i(\pi)$ is the set that captures i 's *belief* or *subjective uncertainty* about other agents' strategies in a strategy profile π , hence we will call it often *belief set*. Yet another description is that instead of writing a belief as a distribution over profiles, the belief of agent i is written as a *point estimate* π_{-i} plus an uncertainty parameter r_i (which together induce a set \mathcal{B}_i). The use of a ball to capture uncertainty is both motivated by its formal simplicity, and the degree of freedom it provides which is set aside from any obvious domain specific/context-dependent constraints.

Notice that as r_i approaches 0, i becomes almost sure about other agents' strategies, and thus \mathcal{B}_i becomes $\{\pi_{-i}\}$. Noteworthy is that for two-player games r reduces to the distance between two probability distributions. For more than two players, any distance on probabilities induces a natural metric d on uncorrelated profiles where for each $j \neq i$ we consider all π'_j close to π_j .

Local responses

Let Π_i denote the set of all strategies available to i . First, we introduce some notions of best response.

Definition 6. π'_i locally dominates π_i in the set $\Pi_{-i}^* \subseteq \Pi_{-i}$ if: (a) for all $\pi_{-i} \in \Pi_{-i}^*$, $u_i(\pi'_i, \pi_{-i}) \geq u_i(\pi_i, \pi_{-i})$; and (b) there exists $\pi'_{-i} \in \Pi_{-i}^*$ such that $u_i(\pi'_i, \pi'_{-i}) > u_i(\pi_i, \pi'_{-i})$. π'_i strictly locally dominates π_i in Π_{-i}^* if (a) holds with strict inequality.

Note that when $\Pi_{-i}^* = \Pi_{-i}$, local dominance and strict local dominance boil down to weak and strict strategic dominance, respectively (Shoham and Leyton-Brown 2008).

Given a mixed strategy profile π , for each $i \in N$ with r_i , a strategy π_i is a distance-based

- **(W)orst-case** best response (or *maximin*) if $\pi_i = \arg \max_{\pi'_i \in \Pi_i} \{\min(u_i(\pi'_i, \pi_{-i}) \mid \pi_{-i} \in \mathcal{B}_i(\pi))\}$.
- **(B)est-case** best response (or *maximax*) if $\pi_i = \arg \max_{\pi'_i \in \Pi_i} \{\max(u_i(\pi'_i, \pi_{-i}) \mid \pi_{-i} \in \mathcal{B}_i(\pi))\}$.
- **(WR) Worst-Case Regret** best response if $\pi_i = \arg \min_{\pi'_i \in \Pi_i} \max\{r \in \mathcal{G}_i(\pi'_i, \pi_{-i}) \mid \pi_{-i} \in \mathcal{B}_i(\pi)\}$ where $r \in \mathcal{G}_i(\pi_i, \pi_{-i}) = \max_{\pi'_i \in \Pi_i} (u_i(\pi'_i, \pi_{-i}) - u_i(\pi_i, \pi_{-i}))$.
- **(U)ndominated** best response if there is no π'_i that locally dominates π_i in the set $\mathcal{B}_i(\pi)$.
- **(D)ominant** best response if π_i locally dominates all π'_i in the set $\mathcal{B}_i(\pi)$.

► **(SD) Strictly dominant** best response if π_i strictly locally dominates all π'_i in the set $\mathcal{B}_i(\pi)$.

By the following proposition, we characterize the relations between these notions. We make no assumption on the metric since it only uses single sets of possible strategy profiles (of opponents), a.k.a. belief sets. In other words, these relations are independent from the choice of the metric.

Theorem 2. *Given any ignorance factor r , the following statements hold:*

- (a) *If π_i is a SD_r best response, then π_i is a D_r best response.*
- (b) *If π_i is a D_r best response, then π_i is a W_r and B_r best response.*
- (c) *If π_i is a D_r best response, then π_i is a WR_r best response.*
- (d) *If π_i is a unique W_r, B_r or WR_r best response, then π_i is a U_r best response.*

Proof. (a) Easily follows by Definition 6.

(b) Consider any action $\pi'_i \neq \pi_i$. Since $u_i(\pi_i, \pi_{-i}) \geq u_i(\pi'_i, \pi_{-i})$ for any state $\pi_{-i} \in \mathcal{B}_i(\pi)$, this holds in particular for the states with maximal utility and minimal utility.

(c) If π_i is a D_r best response, then $\text{reg}_i(\pi_i, \pi_{-i}) = \max_{\pi'_i \in \Pi_i} (u_i(\pi'_i, \pi_{-i}) - u_i(\pi_i, \pi_{-i})) \leq \max_{\pi'_i \in \Pi_i} (u_i(\pi_i, \pi_{-i}) - u_i(\pi_i, \pi_{-i})) = 0$ for all $\pi_{-i} \in \mathcal{B}_i(\pi)$. The regret of any other action π'_i can only be higher, thus π'_i is a WR_r response.

(d) Suppose that π_i is a B_r best response. If π_i is not a U_r response, then there is an action $\pi'_i \neq \pi_i$ that locally dominates π_i . In particular, $u_i(\pi'_i, \pi_{-i}^*) \geq u_i(\pi_i, \pi_{-i}^*)$ in the best state π_{-i}^* , which means that π'_i is also a B_r best response. Note that uniqueness is a necessary condition, otherwise, we can consider two actions π_i, π'_i that have the same utility in the best case, but one of them dominates the other. The proof for W_r and WR_r is similar. \square

The following result also holds for any metric since it only uses containmentment.

Proposition 3. *If $r'_i < r_i$ then $SD_{r'_i}$ response implies $SD_{r'_i}$ response.*

Proof. Since $\mathcal{B}(\pi, r'_i) \subseteq \mathcal{B}(\pi, r_i)$, condition (a) of Def. 6 must hold in all states. \square

This does not hold in any of the other variations. To see this, consider W -equilibrium, since for any response π_i , $\min\{u_i(\pi_i, \pi_{-i}) \mid \pi_{-i} \in \mathcal{B}_i(\pi, r)\} \leq \min\{u_i(\pi_i, \pi_{-i} \mid \mathcal{B}_i(\pi, r^*))\}$ whenever $\mathcal{B}_i(\pi, r^*) \subseteq \mathcal{B}_i(\pi, r)$. The other variations are similar.

Equilibrium

Now, we are ready to define the notion of *distance-based* equilibrium. Let $\mathbf{r} := (r_1, \dots, r_n)$ be the *ignorance vector* which stores ignorance factor for each agent $i \in N$. Assume that $\star \in \{W, B, WR, U, D, SD\}$. Then, π is called a **distance-based \star_r -equilibrium** if for every agent i , whose belief set $\mathcal{B}_i(\pi)$ is defined w.r.t. r_i where $r_i = \mathbf{r}_i$, π_i is a \star -best response. When all the agents have the same r , we




Figure 1: Entailment of equilibrium concepts; An arrow to \star_r means that there is some $r > 0$ for which the entailment holds. An arrow from \star_r means the entailment holds for any $r > 0$ (except for the dotted arrow which is slightly weaker, see Prop. 8). Dashed arrows mark entailments that are known or obvious.

will use r instead of \mathbf{r} as a subscript, or totally omit it whenever it is clear from the context. Similar solution concepts to U_r, W_r, WR_r (in specific games) are studied in works (Meir, Lev, and Rosenschein 2014; Meir and Parkes 2015).

Observation 4. *For all definitions above, if $r_i = 0$ for all $i \in N$, then a \star -equilibrium is a Nash equilibrium.*

The statements below follow immediately from the relations between corresponding responses (i.e., Theorem 2).

Corollary 5. *Given any r , the following statements hold:*

- (a) *If π is a SD_r -equilibrium, then π is a D_r -equilibrium.*
- (b) *If π is a D_r -equilibrium, then π is a W_r, B_r and WR_r -equilibrium.*
- (c) *If π is a W_r, B_r or WR_r -equilibrium, then π is a U_r -equilibrium.*

A brief summary of our results is illustrated in Figure 1.

Locally Best Response and Trembling Hand

Definitions 1-4 capture stability of equilibrium in a rigorous formal way, but require reasoning about sequences of profiles that do not seem have a clear cognitive interpretation.

In this section we aim to get a better understanding of these concepts and of our distance-based equilibrium concepts, by exploring the connections between them. Moreover, the proposed distance-based best-response (and hence equilibrium) do have cognitive interpretation which is the observational limitation that an agent has.

We first argue that robustness (under the appropriate Def. 1-4) implies stability under uncertainty (under Def. 5) when the ignorance factors of all agents are sufficiently small.

Observation 6. *Given a profile π and an agent $i \in N$, for any mixed strategy π_i , $\max_{\pi_{-i} \in \mathcal{B}_i(\pi)} u_i(\pi_i, \pi_{-i}) = \min_{\pi_{-i} \in \mathcal{B}_i(\pi)} u_i(\pi_i, \pi_{-i}) = u_i(\pi_i, \pi_{-i})$ if $r_i = 0$.*

Having Observation 6 in mind, realize that the classical notion of MN implies our notions of W_r, B_r, WR_r and U_r equilibria for $r = 0$, but not for any $r > 0$, as there can be (even pure) Nash equilibria that are weakly dominated.

The following result provides a partial picture.

Proposition 7. *Given a normal-form game \mathcal{G} , if π is a strict Trembling-hand perfect equilibrium, then there is an $\epsilon > 0$ such that if $r_i < \epsilon$ for every $i \in N$, then π is a U_r -equilibrium.*

Proof. Assume a game \mathcal{G} and a Trembling-hand perfect equilibrium π in \mathcal{G} . By Definition 1, there is a sequence $\{\pi^k\}_{k=0}^\infty$ of totally mixed strategies which converges to π . Take an arbitrary $\epsilon > 0$. By convergence there is a K such that $d(\pi^k, \pi) < \epsilon$ whenever $k \geq K$. Moreover, it follows that for every $i \in N$, $d(\pi_{-i}^k, \pi_{-i}) < \epsilon$ as well for $k \geq K$. For each agent $i \in N$, let $r_i = d(\pi_{-i}^K, \pi_{-i})$ i.e., $\mathbf{r} = (r_1, \dots, r_n)$. Now, for any agent i , we know that $\pi_{-i}^k \in \mathcal{B}_i(\pi)$ whenever $k \geq K$, and also since π is strict- T , π_i is a best response to π_{-i}^k . This shows that π_i is U_{r_i} -response and that π is a U_r equilibrium. \square

In the other direction, it seems that a D_r -equilibrium (w.r.t. any $r > 0$) must be a strict Trembling-hand perfect equilibrium.

Proposition 8. *If there is an r^* , such that π is D_r -equilibrium for all $r \in (0, r^*)$, then π is strict- T .*

Proof. Consider an arbitrary sequence of $r^k < r^*$ that converges to 0. Since π is a D_{r^k} -equilibrium, then for every player $i \in N$, π_i is a best response to every $\pi_{-i} \in \mathcal{B}(\pi, r^k)$, and a strict best response to at least one profile $\pi_{-i}^{k,i}$. Moreover, $\pi_{-i}^{k,i}$ is w.l.o.g. totally mixed (we can mix it with a low probability for any other profile such that π_i remains a best-response). Let $\pi_i^{r^k}$ be a mixed strategy that selects π_i with probability $1 - r^k$.

We therefore get n sequences of totally-mixed profiles converging to π , where in each sequence $(\langle \pi^{k,i}, \pi_i^{r^k} \rangle)_{k=0}^\infty$, π_i is a strict best-response to the entire sequence.

Let $\pi^k = (\pi^{k,i}, \pi_i)$ for all k such that $k \bmod n = (i-1)$ (that is, we interleave subsequences). Note that $(\pi^k)_{k=0}^\infty$ converges to π , and for every agent i there is a subsequence for which π_i is a strict best-response. Thus π is a strict Trembling-hand perfect equilibrium (strict- T). \square

Proposition 9. *For any r , if π is D_r -equilibrium, then π is TP . Similarly, SD_r entails $s-TP$.*

Proof. Fix any r . It follows from the fact that whatever sequence we choose in $\mathcal{B}_i(\pi_{-i}, r)$, π_i will be a best response [respectively, strict best response] to it by definition, hence satisfying the condition of TP [$s-TP$]. \square

Proposition 10. *If π is a TP , then there is an r such that π is a U_r -equilibrium. Further, if the game also is generic then π is a D_r -equilibrium.*

Proof. Assume that π is a TP , then for each sequence $\{\pi^k\}_{k=0}^\infty$ of totally mixed strategy profiles which converge to π , there is a K such that π_i is a best response to π_{-i}^k for all $k \geq K$. All possible sequences form a ball for each player, and we take the infimum of K of those (all) possible sequences; call it K^* . Then, for all $i \in N$, we define $r_i = d(\pi_{-i}, \pi_{-i}^{K^*})$. Obviously, either π_i locally dominates all the other responses to every $\pi_{-i} \in \mathcal{B}_i(\pi, r_i)$ for every player i (which implies that π is a D_r -equilibrium), or there is some response that has the same utility as π_i (which means non-genericity). \square

The following result provides a link between strict- TP and SD_r -equilibrium.

Proposition 11. *If π is a strict- TP then there is an $\epsilon > 0$ such that if $r_i < \epsilon$ for all $i \in N$ then π is SD_r -equilibrium.*

Proof. Assume that π is not a SD_r -equilibrium for any $r > 0$. We will construct a sequence of states $\{\pi^k\}_{k=0}^\infty$ that converges to π , but such that for every k there is some agent i for which π_i is not a strict best response to π_{-i}^k . Let $r_k = \frac{1}{k}$. Since π is not a SD_r -equilibrium, there is some $i \in N$, a profile $\pi_{-i}^k \in \mathcal{B}_i(\pi, r_k)$, and an action a'_i such that $u_i(\pi_{-i}^k, \pi_i) \leq u_i(\pi_{-i}^k, a'_i)$. That is, π_i is not a strict best response to π_{-i}^k . We set $\pi^k = (\pi_{-i}^k, \pi_i)$. By construction, $d(\pi^k, \pi) \leq \frac{1}{k}$ and thus $\{\pi^k\}_{k=0}^\infty$ converges to π . \square

It seems $s-T$ -equilibrium does not imply D_r -equilibrium. The proof is included in the longer version of the paper due to space restrictions.

Proposition 12. *There is a $s-T$ -equilibrium that is not a B_r equilibrium for any $r > 0$ (and thus not a D_r -equilibrium).*

Proposition 13. *Given a game \mathcal{G} , if π is an R -equilibrium, then there is an $r > 0$ such that π is a B_r and W_r -equilibrium.*

Proof. We give a proof for B_r (W_r is similar) case. Assume that π is an R , then by definition 5, there is an ϵ such that for every $\epsilon^* \leq \epsilon$, π^* such that $\pi^* \rightarrow \pi$ is an equilibrium. Let π' be a profile such that $d(\pi', \pi) = 0$. By robustness, π' is an equilibrium. Now, for every player $i \in N$, observe that $d(\pi', \pi) > d(\langle \pi_i, \pi'_{-i} \rangle, \pi)$. Then, $\langle \pi_i, \pi'_{-i} \rangle \in \mathcal{B}_i(\pi', \epsilon)$. Hence, π_i is a best response to any sequence of $\{\langle \pi_i, \pi'_{-i} \rangle\}$ converging to π . Therefore, π_i must be a B_r best response for $r = d(\langle \pi_i, \pi'_{-i} \rangle, \pi)$. \square

Proposition 14. *Let a be a pure profile in D_r -equilibrium, then a is $\frac{r}{n^{1/2}}$ -Robust. Moreover, any ϵ -Robust equilibrium is a D_ϵ -equilibrium.*

Proof. Assume that $a \in D_r$ for some r , and consider some ϵ -noisy variant of a . $d(a_{-i}, \pi_{-i}) = \left(\sum_{j \neq i} (1 - \pi(a_j))^2 \right)^{1/2} \leq (n\epsilon^2)^{1/2}$, thus $\pi_{-i} \in \mathcal{B}_i(a, r)$ for $r = (n\epsilon^2)^{1/2} = n^{1/2}\epsilon$. This means that a is ϵ -robust for $\epsilon = \frac{r}{n^{1/2}}$. In the other direction, suppose that a is ϵ -Robust and consider some π . If $\pi(a_j) > 1 - \epsilon$, then $d(\pi, a) \geq ((\pi(a_j) - 1)^2)^{1/2} > \epsilon$, which means that all

vectors in $\mathcal{B}_i(a, r)$ are (at most) r -noisy variants of a . Thus for any $r \leq \epsilon$, a_i is a best-response to all of $\mathcal{B}_i(a, r)$, and is therefore a D_ϵ equilibrium. \square

The following result follows immediately from the Proposition 13 and Corollary 5.

Corollary 15. *Given a game \mathcal{G} , if π is an R , then there is an $r > 0$ such that π is a U_r -equilibrium.*

See Figure 1, for a summary of the results obtained in this Section.

Existence Results

In the previous section, we had noted that MN implies our distance-based notions of W_r , WR_r , B_r and U_r equilibria as a special case when $r = 0$. In this section, we deliver existence results regarding the several variants of distance-based equilibria.

An immediate negative result to start with, D_r -equilibrium which is central in Figure 1, seems to be too strong to exist in general. To see this, assume a game in which all actions have the same payoffs for every player. Obviously any profile is a Nash equilibrium, yet none of them is a D_r -equilibrium (including $r = 0$) due to (b) of Definition 6 of local dominance.

Corollary 16. *D_r -equilibrium does not exist in general.*

Balancing out this negative news, the remaining distance-based equilibria entailed by D_r do exist.

Theorem 17. *Every finite normal-form game \mathcal{G} has a W_r , B_r and WR_r -equilibria.*

We omit the actual long proof here due to space limitations, and present it in the extended (ArXiv) version. The proof idea is based on well-known application of Kakutani's fixed point's theorem i.e., defining the best response correspondence (for each one of them), and showing its convexity and upper-continuity. It is rather straightforward in the cases of W_r and B_r -equilibria. In the case of WR_r -equilibrium, it is obtained by showing the convexity and piece-wise linearity of the worst-case best response expression (6).

Next, the following result immediately follows from the above theorem and Corollary 5.

Corollary 18. *Every finite normal-form game \mathcal{G} has a U_r -equilibrium.*

Discussion through Examples

To provide a better intuition, we give examples of well-known 2×2 games from the basic game theory literature, and compare the outcomes of standard notions of equilibria against some notions of equilibria that we defined via local responses.

For convenience, the assumed metric is Euclidean; hence, if the opponent plays a mixed strategy $(x, 1 - x)$, the player believes that the strategy is anywhere in the set $\{(y, 1 - y) : \max\{0, x - r\} \leq y \leq \max\{1, x + r\}\}$.

	Left	Right
Up	1, 1	2, 0
Down	0, 2	2, 2

Figure 2: Trembling-Hand Game where both $\langle \text{Up, Left} \rangle$ and $\langle \text{Down, Right} \rangle$ are PN , yet only $\langle \text{Up, Left} \rangle$ is T .

Trembling-Hand Game

Call the example given in Figure 2 Trembling-Hand Game for demonstration purposes. It seems that the pure strategy Nash equilibrium $PN = \{(\text{Up, Left}), (\text{Down, Right})\}$ while Trembling-hand perfect equilibrium $T = \{(\text{Up, Left})\}$ which matches with W .

Assume that $r_1 = r_2 \approx 0.14$. Now consider two mixed Nash strategy equilibria $\pi = \langle (1, 0), (1, 0) \rangle$ and $\pi' = \langle (0, 1), (0, 1) \rangle$ where the former is also a T . See that for $\mathcal{B}_1(\pi)$, every strategy is dominated by π_1 in terms of B_r -best response and W_r -best response (analogous for the second agent). Moreover, the regret increases as agent 1 diverges from $(1, 0)$. Hence $\pi \in B_r \cap W_r$.⁵ In the case of π' , it is a B_r -response since payoffs are already 2 for both agents. On the other hand, $\pi' \notin W_r$ since worst case keeps improving for any agent who keeps deviating. Therefore, the regret also gets minimized since the best case value is fixed at 2.

Matching Pennies

In the game of *Matching Pennies*, $PN = \emptyset$ whereas The set of mixed strategy Nash equilibria is a singleton i.e., $MN = \{\pi\}$ where $\langle (0.5, 0.5), (0.5, 0.5) \rangle$. If $r \geq 0.5$ then

	Heads	Tails
Heads	1, -1	-1, 1
Tails	-1, 1	1, -1

Figure 3: Matching Pennies

all profiles are B_r -equilibrium. In the case of U_r , any strategy in the set $(0.5 - r, 0.5 + r)$ (coupled symmetrically with the other player) forms an equilibrium. This means an imprecise randomization can also be an equilibrium, provided that players have some level of uncertainty over the exact randomization of the other player.

Stag Hunt

	Stag	Hare
Stag	5, 5	-1, 3
Hare	3, -1	1, 1

Figure 4: Stag-Hunt Game

We continue with the rather well-known coordination game Stag-Hunt illustrated in Figure 4.

⁵Indeed, for such a value of r , the strategy of the opponent varies only by 0.1. And the worst case is defined by the case that the opponent plays $(1, 0)$.

Pure strategy-Nash equilibria are $S = \langle \text{Stag, Stag} \rangle$ and $H = \langle \text{Hare, Hare} \rangle$. There is also a mixed equilibrium that we will not consider. Playing S is more socially desirable, but less stable according to several solution concepts based on uncertainty and risk aversion (see the work by (Carlsson and Van Damme 1993) for an overview and discussion). However these notions are not sensitive to possible differences between players' perceptions and do not quantify the instability of S . In contrast, such quantification is very natural when we consider distance-based equilibria, by the size r of the respective belief set. For example, for pessimistic players (who attribute excessive probability to the opponent playing "Hare"), S is a W_r -equilibrium for any r , but H is only a W_r -equilibrium if $r \leq \frac{1}{3}$.

A Bliss of Ignorance

Once a new notion of equilibrium is introduced, it is natural to ask whether it leads to any efficient outcome in the game. In this section, we explore this question and provide with an exemplary class of normal-form games that this is indeed the case. In doing so, we employ the notion of *Price of Anarchy* which is central as a measure of how much a system becomes inefficient due to selfish behaviour (Nisan et al. 2007) (recall the preliminaries section for formal definition). The class of games we introduce is a consensus game (Balcan, Blum, and Mansour 2009) which is asymmetric in payoffs.

Definition 7 (Consensus Game). *A normal-form game $\mathcal{G} = (N, A, u)$ where $A_i \geq 2$ for every $i \in N$ is called a consensus game if $u_i(a') = c'$ and $u(a) = c$ for every $a \in A \setminus \{a'\}$ with $c' > c$.*

Intuitively, the given consensus game is a coordination game in which only a single pure strategy profile has a higher payoff i.e., c' for every player compared to all the other pure strategy profiles which has c . Such game model (group) scenarios in which every member player has to agree unilaterally a decision to be taken (e.g., World Trade Org.).

One can observe that any given consensus game has at least two pure strategy equilibria a' and a such that only one of them has a more desirable outcome i.e., $SW(a') = n \cdot c'$, and $SW(a) = n \cdot c$. The following result shows that an ignorance factor $r > 0$, eliminates the undesirable equilibrium.

Proposition 19. *Every consensus game has a unique D_r equilibrium where $r_i > 0$ for all $i \in N$. Moreover, PoA is 1.*

Proof. Observe that profile a' has the best possible payoff for every single agent, hence it is an equilibrium. Moreover, due to the linearity of utilities (i.e., $\pi_i(a')c' + (1 - \pi_i(a'))c > \pi'_i(a')c' + 1 - \pi'_i(a')c$ whenever $\pi_i(a') > \pi'_i(a')$), pure strategy $\pi_i(a') = 1$ locally dominates every other strategy (i.e., $\alpha = 1$) for any $\mathcal{B}_i(\pi'', r)$ with $r > 0$, hence it is unique (since, by Definition 6 there cannot be two distinct best response which can locally dominate each other). As it is the best possible outcome, PoA becomes 1. \square

Exploring such scenarios and extending them to more general class of games is left as future work. Yet still to develop a general understanding, it is important to look at PoA from the lenses of distance-based uncertainty. In this regard,

we deliver our final technical result. In particular, we provide a bound (in terms of r) on the gain/loss of social welfare in an equilibrium modulo strict uncertainty.

The smoothness framework provides a convenient tool to bound the PoA in games (Roughgarden 2009): If there are $\lambda, \mu > 0$ s.t. for any two pure profiles a, a' we have

$$\sum_{i \in N} u_i(a'_i, a_{-i}) \geq \lambda \sum_{i \in N} u_i(a') - \mu \sum_{i \in N} u_i(a),$$

then for any pure/mixed/correlated/coarse-correlated equilibrium π^* and any profile \vec{a} :

$$\frac{SW(\pi^*)}{SW(a')} \geq \frac{\lambda}{1 + \mu}.$$

The proof is trivial for pure equilibria. Now, the question we ask is "can we extend this result to \star_r equilibria (perhaps with a relaxed bound)?"

For a game \mathcal{G} , let

$$\delta_G(r) = \max\left\{\max\left\{\frac{u_i(a_i, \pi_{-i})}{u_i(a_i, \pi'_{-i})}, \frac{u_i(a_i, \pi'_{-i})}{u_i(a_i, \pi_{-i})}\right\} : i \in N, a_i \in A_i, \pi'_{-i} \in \mathcal{B}_i(\pi_{-i}, r)\right\}, \quad (1)$$

i.e., the maximal utility ratio of an agent within a sphere of radius r .

Theorem 20. *If there are $\lambda, \mu > 0$ s.t. for any two pure profiles a, a' we have*

$$\sum_{i \in N} u_i(a'_i, a_{-i}) \geq \lambda \sum_{i \in N} u_i(a') - \mu \sum_{i \in N} u_i(a),$$

then for any \star_r -pure equilibrium a^* (for any $\star \in \{U, D, W, B\}$) and any profile a' :

$$\frac{SW(a^*)}{SW(a')} \geq \frac{\lambda}{\delta_G(r)^2 + \mu}.$$

Due to space limitations, we move the proof to the appendix.

Conclusion and Future Avenues

We have introduced a distribution-free agent model based on strict uncertainty, and studied consequent equilibria notions under different best-response behaviours. In the context of normal-form games, we explored the links between the notions we defined and a handful of existing well-known solution concepts which model mistakes and imprecision such as Trembling-hand perfect equilibrium (variants) and Robust equilibrium. For instance, it is shown that our notion is naturally generalizes Robust equilibrium. It seems that strict equilibrium notion D_r does not exist in general while all other entailed distance-based notions exist. Complementing those existence results with complexity results is an interesting line of future work.

We looked for a possible scenario in which such solution concepts could potentially be useful, and introduced a coordination game in which ignorance was indeed helpful for the players to avoid a worst-outcome. Investigating more general game classes that distance-based uncertainty solutions

give rise to nice outcome guarantees deserves a further study on its own, and is our high priority for future research. As a more general outlook, we showed how to bound the loss of social welfare in any equilibrium (PoA) as uncertainty grows in terms of ignorance factor r . It would be nice to obtain finer bounds for games with different local-best responses.

Moreover, studying these notions on certain classes of games e.g., repeated games, as well as extending to extensive form games in general is our future research agenda.

References

- Aghassi, M., and Bertsimas, D. 2006. Robust game theory. *Mathematical Programming* 107(1-2):231–273.
- Apt, K. R. 2007. The many faces of rationalizability. *The BE Journal of Theoretical Economics* 7(1).
- Aumann, R. J. 1997. Rationality and bounded rationality. *Games and Economic Behavior* 21(1-2):2–14.
- Balcan, M.-F.; Blum, A.; and Mansour, Y. 2009. Improved equilibria via public service advertising. In *ACM-SIAM'09*, 728–737. Society for Industrial and Applied Mathematics.
- Camerer, C. F.; Ho, T.-H.; and Chong, J.-K. 2004. A cognitive hierarchy model of games. *The Quarterly Journal of Economics* 119(3):861–898.
- Caragiannis, I.; Kurokawa, D.; and Procaccia, A. D. 2014. Biased games. In *Proc. of AAAI'14*, 609–615.
- Carlsson, H., and Van Damme, E. 1993. Global games and equilibrium selection. *Econometrica: Journal of the Econometric Society* 989–1018.
- Cohen, E. R. 1998. *An introduction to error analysis: The study of uncertainties in physical measurements*. IOP Publishing.
- Conitzer, V.; Walsh, T.; and Xia, L. 2011. Dominating manipulations in voting with partial information. In *Proc. of AAAI'11*, 638–643.
- Deligkas, A.; Fearnley, J.; and Spirakis, P. 2016. Lipschitz continuity and approximate equilibria. In *International Symposium on Algorithmic Game Theory*, 15–26. Springer.
- Dow, J., and Werlang, S. R. d. C. 1994. Nash equilibrium under knightian uncertainty: breaking down backward induction. *Journal of Economic Theory* 64(2):305–324.
- Farina, G.; Marchesi, A.; Kroer, C.; Gatti, N.; and Sandholm, T. 2018. Trembling-hand perfection in extensive-form games with commitment. In *Proc. of IJCAI'18*, 233–239.
- Farina, G.; Gatti, N.; and Sandholm, T. 2018. Practical exact algorithm for trembling-hand equilibrium refinements in games. In *Proc. of NIPS'17*, 5044–5054.
- Fudenberg, D., and Tirole, J. 1991. *Game Theory*. MIT Press.
- Gilboa, I., and Schmeidler, D. 1989. Maxmin expected utility with non-unique prior. *Journal of mathematical economics* 18(2):141–153.
- Gilboa, I.; Postlewaite, A. W.; and Schmeidler, D. 2008. Probability and uncertainty in economic modeling. *Journal of Economic Perspectives* 22(3):173–88.
- Halpern, J. Y. 2017. *Reasoning about uncertainty*. MIT press.
- Heckerman, D.; Mamdani, A.; and Wellman, M. P. 1995. Real-world applications of bayesian networks. *Communications of the ACM* 38(3):24–26.
- Hyafil, N., and Boutilier, C. 2004. Regret minimizing equilibria and mechanisms for games with strict type uncertainty. In *In Proc. of UAI'04*, 268–277. AUAI Press.
- Kohlberg, E. 1981. Some problems with the concept of perfect equilibrium. In *NBER Conference on Theory of General Economic Equilibrium*.
- Lev, O.; Meir, R.; Obraztsova, S.; and Polukarov, M. 2019. Heuristic voting as ordinal dominance strategies. In *Proc. of AAAI'19*, 2077–2084.
- Marinacci, M. 2000. Ambiguous games. *Games and Economic Behavior* 31(2):191–219.
- McKelvey, R. D., and Palfrey, T. R. 1995. Quantal response equilibria for normal form games. *Games and economic behavior* 10(1):6–38.
- Meir, R., and Parkes, D. 2015. Congestion games with distance-based strict uncertainty. In *Proc. of AAAI'15*, 986–992.
- Meir, R.; Lev, O.; and Rosenschein, J. S. 2014. A local-dominance theory of voting equilibria. In *Proc. of EC '14*, 313–330.
- Messner, M., and Polborn, M. 2005. Robust political equilibria under plurality and runoff rule. *IGIER Working Paper*.
- Nisan, N.; Roughgarden, T.; Tardos, E.; and Vazirani, V. V. 2007. *Algorithmic game theory*. Cambridge University Press.
- Okada, A. 1981. On stability of perfect equilibrium points. *International Journal of Game Theory* 10(2):67–73.
- Potyka, N.; Acar, E.; Thimm, M.; and Stuckenschmidt, H. 2016. Group decision making via probabilistic belief merging. In *Proc. of IJCAI'16*, 3623–3629.
- Roughgarden, T. 2009. Intrinsic robustness of the price of anarchy. In *Proc. of STOC'09*, 513–522.
- Savage, L. J. 1951. The theory of statistical decision. *Journal of the American Statistical Association* 46(253):55–67.
- Selten, R. 1975. Reexamination of the perfectness concept for equilibrium points in extensive games. *International journal of game theory* 4(1):25–55.
- Shoham, Y., and Leyton-Brown, K. 2008. *Multiagent systems: Algorithmic, game-theoretic, and logical foundations*. Cambridge University Press.
- Thrun, S.; Fox, D.; Burgard, W.; and Dellaert, F. 2001. Robust Monte Carlo localization for mobile robots. *Artificial intelligence* 128(1-2):99–141.
- Wald, A. 1939. Contributions to the theory of statistical estimation and testing hypotheses. *The Annals of Mathematical Statistics* 10(4):299–326.
- Williamson, T. 1992. Inexact knowledge. *Mind* 101(402).